

UNIVERSITY OF CALIFORNIA
Santa Barbara

Medium Access Control and Bandwidth
Allocation in Millimeter-Wave and WiFi
Networks

A Dissertation submitted in partial satisfaction
of the requirements for the degree of

Doctor of Philosophy

in

Electrical and Computer Engineering

by

Sumit Singh

Committee in Charge:

Professor Upamanyu Madhow, Co-Chair

Professor Elizabeth M. Belding, Co-Chair

Professor Jerry Gibson

Professor João Hespanha

December 2009

The Dissertation of
Sumit Singh is approved:

Professor Jerry Gibson

Professor João Hespanha

Professor Elizabeth M. Belding, Committee Co-Chairperson

Professor Upamanyu Madhow, Committee Co-Chairperson

November 2009

Medium Access Control and Bandwidth Allocation in Millimeter-Wave and
WiFi Networks

Copyright © 2009

by

Sumit Singh

To Ma, Papa, Swati, Sweta.

Acknowledgments

I am indebted to my advisors, Professor Upamanyu Madhow and Professor Elizabeth Belding, for a wonderful educational experience over the last five years. Their advice, attention, and research insights, along with their patience and help with my writing have made this dissertation possible. Their professionalism and passion for research inspire me.

I thank Professor Jerry Gibson and Professor João Hespanha for serving on my thesis committee. Thanks to Dr. Janet Kayfetz for her whole-hearted help with my writing and presentation skills.

I thank my collaborators Raghu Mudumbai, Prashanth Acharya and Federico Ziliotto for their contributions to this dissertation. I have learned a lot from the many stimulating discussions with them. I am grateful to Professor Mark Rodwell and Munkyo Seo for the collaborative work on 60 GHz networking.

I thank my fellow lab members from both WCSL and MOMENT labs for their helpful comments and suggestions on my research, paper drafts, and practice talks. Their friendly company, and insightful and/or entertaining discussions made for a great work environment.

I thank Dr. Murali Kodialam and Dr. T. V. Lakshman for mentoring me during my internship at Bell Labs in 2007, and Dr. Lily Yang and Dr. Guoqing

Li for my internship at Intel Communication Technology Labs in 2008. Both my internships were great learning experiences.

I thank my undergraduate advisor at IIT Bombay, Professor D. Manjunath, for introducing me to the field of networking and for his help and encouragement for my graduate education.

The past five years at UCSB have been very enriching and fun because of the intellectually stimulating environment, the wide range of interesting courses offered by the UCSB faculty, the extracurricular activities and events at school, and the beautiful campus. I consider myself fortunate to have been a graduate student at UCSB.

Thanks to all my friends for their refreshing company that has kept my spirits high, and for their help over the years. The incredible support of my closest friends has always been a treasure. Words cannot express my gratitude towards my family for their love and encouragement throughout my life. This dissertation is dedicated to them.

Curriculum Vitæ

Sumit Singh

Education

- 2009 Ph.D. in Electrical and Computer Engineering,
University of California, Santa Barbara, CA, USA.
Advisors: Prof. Upamanyu Madhow and Prof. Elizabeth Belding
- 2008 Certificate - Graduate Program in Management Practice,
University of California, Santa Barbara, CA, USA.
- 2006 Master of Science in Electrical and Computer Engineering,
University of California, Santa Barbara, CA, USA.
- 2002 Bachelor of Technology in Electrical Engineering,
Indian Institute of Technology Bombay, Mumbai, India.

Experience

- 2006 – 2009 Research Assistant, University of California, Santa Barbara, CA.
- Summer 2008 Intern, Intel Communication Technology Labs, Hillsboro, OR.
- Summer 2007 Intern, Bell Labs, Alcatel-Lucent, Murray Hill, NJ.
- 2004 – 2005 Teaching Assistant, ECE Department, University of California,
Santa Barbara, CA.
- 2002 – 2004 Engineer, Samsung India Software Operations, Bangalore, India.

Publications

- S. Singh, F. Ziliotto, U. Madhow, E. M. Belding, and M. Rodwell. **Blockage and Directivity in 60 GHz Wireless Personal Area Networks: From Cross-Layer Model to Multihop MAC Design.** IEEE Journal on Selected Areas in Communication (JSAC): Special Issue on Realizing Gbps Wireless Personal Area Networks, vol 27, no. 8, pp. 1400-1413, Oct. 2009.
- S. Singh, P. Acharya, U. Madhow, and E. M. Belding. **Sticky CSMA/CA: Implicit Synchronization and Real-time QoS in Mesh Networks.** Ad hoc Networks Journal, vol.5, no. 6, pp. 744-768, Aug. 2007.
- S. Singh, U. Madhow, and E. M. Belding. **Shaping Throughput Profiles in Multihop Wireless Networks: A Resource Biasing Approach.** Submitted.
- S. Singh, R. Mudumbai, and U. Madhow. **Distributed Coordination with Deaf Neighbors: Efficient Medium Access for 60 GHz Mesh Networks.** Submitted.
- R. Mudumbai, S. Singh, and U. Madhow. **Medium Access Control for 60 GHz Outdoor Mesh Networks with Highly Directional Links.** in Proc. IEEE INFOCOM 2009, Mini Conference, Rio de Janeiro, Brazil, Apr. 2009.
- S. Singh, U. Madhow, and E. M. Belding. **Beyond Proportional Fairness: A Resource Biasing Framework for Shaping Throughput Profiles in Multihop Wireless Networks.** in Proc. IEEE INFOCOM 2008, Mini Conference, Phoenix, AZ, Apr. 2008.

- S. Singh, F. Ziliotto, U. Madhow, E. Belding, and M. Rodwell. **Millimeter Wave WPAN: Cross-Layer Modeling and Multihop Architecture.** in Proc. IEEE INFOCOM 2007, Minisymposium, Anchorage, AL, May 2007.
- P. Acharya, S. Singh, and H. Zheng. **Reliable Open Spectrum Communications Through Proactive Spectrum Access.** in Proc. International Workshop on Technology and Policy for Accessing Spectrum (TAPAS), Boston, MA, Aug. 2006.
- P. Manohar, A. Padmanath, S. Singh, and D. Manjunath. **Multiperiod Virtual Topology Design in Wavelength Routed Optical Networks.** IEE Proc. - Circuits, Devices and Systems, vol 150, no. 6, pp. 516-520, Dec. 2003.

Abstract

Medium Access Control and Bandwidth Allocation in Millimeter-Wave and WiFi Networks

Sumit Singh

The past decade has witnessed a wireless revolution with the proliferation of wireless devices, an explosive increase in the number of users, and a surge in bandwidth-hungry multimedia applications over wireless networks. Current generation wireless networks are not equipped to handle the stringent demands on network resources imposed by the confluence of these wireless usage trends. Motivated by the imminent need for effective solutions to this problem, we explore a few promising approaches to meet the increasing demand for higher data rates. Our primary focus is on two fundamental components of a network architecture: wireless medium access control (MAC) and bandwidth allocation.

The first part of this dissertation investigates medium access control for a new generation of short-range indoor and outdoor wireless networks with multiGigabit link speeds, envisioned to operate in the unlicensed 60 GHz “millimeter (mm) wave” band. In addition to the abundant unused bandwidth (e.g., 57-64 GHz in the U.S.), and the potential for high spatial reuse, we are motivated by the recent advances in low-cost mm wave transceiver designs using inexpensive silicon processes. In the second part of this dissertation, we rethink MAC design for efficient

support of real-time and delay-insensitive traffic over mesh networks operating in the 2.4 and 5.8 GHz WiFi bands. Specifically, we present Sticky CSMA/CA, a distributed MAC protocol that achieves time division multiplexing-like performance for real-time flows such as voice over IP (VoIP), without requiring network-wide synchronization. Finally, the third part of this dissertation presents a flexible framework for sharing the transport capacity of multihop wireless networks so as to achieve a desired tradeoff between throughputs achieved by flows and their resource usage, while ensuring efficient network utilization. We show that our resource sharing mechanisms can be implemented as decentralized protocols.

Some key ideas presented in this dissertation have broad applications in wireless network protocol design. Both the Memory-guided Directional MAC (MD-MAC) designed for outdoor 60 GHz mesh networks, and Sticky CSMA/CA, demonstrate the promise of currently under-utilized tools of memory and learning in achieving distributed, implicit network coordination with little overhead. Our 60 GHz MAC protocols are among the few that address networks in which every link is constrained to be directional, in contrast with many prior *pseudo-directional* networking proposals that employ some form of omnidirectional communication for control. Finally, this dissertation demonstrates the crucial role of cross-layer considerations in wireless network protocol design.

Contents

Acknowledgments	v
Curriculum Vitæ	vii
Abstract	x
List of Figures	xv
List of Tables	xviii
1 Introduction	1
1.1 Dissertation Overview	5
1.2 Contributions and Impact	8
1.3 Dissertation Outline	15
2 60 GHz Wireless Personal Area Networks	17
2.1 Introduction	17
2.2 Related Work	22
2.3 Physical Layer Model	25
2.4 Directional MAC Design	39
2.4.1 Discovery Algorithm	40
2.4.2 Normal Mode of Operation	41
2.4.3 Trailing Control Phase	44
2.4.4 Lost Node Discovery and Establishing a Relay Path	45
2.4.5 Achievable Rates	46
2.5 Performance Evaluation	51
2.5.1 Simulation Model	51
2.5.2 Simulation Results	53

2.6	Summary	59
3	60 GHz Outdoor Mesh Networks	60
3.1	Introduction	60
3.2	Related Work	65
3.3	60 GHz Outdoor Mesh Architecture	67
3.4	Network Model	71
3.4.1	Directional Antenna Model	72
3.5	Interference Analysis	75
3.5.1	Protocol Model with Ideal Flat-Top Antennas	76
3.5.2	Protocol Model with General Directional Antennas	78
3.5.3	Physical Model	81
3.5.4	A Pseudo-Wired Abstraction	83
3.6	Simulations	85
3.7	Memory-guided Directional MAC	87
3.7.1	Use of Memory for Implicit Coordination	88
3.7.2	Approximate Protocol Modeling	91
3.7.3	Fairness: Explicit State Reset (ESR)	99
3.7.4	MDMAC Protocol Design Details	100
3.8	Medium Utilization with MDMAC	102
3.9	Simulation Evaluation	105
3.10	Summary	111
4	Sticky CSMA/CA	113
4.1	Introduction	113
4.2	Related Work	118
4.3	The Sticky CSMA/CA Protocol	122
4.3.1	Conventional CSMA/CA	123
4.3.2	Introducing History into CSMA/CA	124
4.3.3	Sticky CSMA/CA	126
4.4	Performance Analysis for a Clique	139
4.4.1	IEEE 802.11b DCF Performance	140
4.4.2	IEEE 802.11b VoIP Call Carrying Capacity	144
4.4.3	Sticky CSMA/CA VoIP Call Carrying Capacity	147
4.5	Performance Evaluation	151
4.5.1	Simulation Setup	152
4.5.2	Line Topology	156
4.5.3	Grid Topologies	162
4.5.4	QoS Performance: Delay and Delay Jitter	168
4.5.5	Effect of Increasing Number of Hops	171

4.5.6	Silence Suppression	172
4.6	Summary	173
5	Shaping Throughput Profiles in Multihop Wireless Networks	178
5.1	Introduction	178
5.2	Related Work	184
5.3	Resource Allocation Framework	188
5.3.1	Computations for a One-Dimensional Network	189
5.3.2	Two-Dimensional Networks	194
5.4	Application to an IEEE 802.11 Multihop Network	200
5.5	Mapping to Network Utility Maximization	206
5.6	Summary	209
6	Conclusions and Future Work	211
6.1	The Big Picture	212
6.2	Future work	214
	Bibliography	219

List of Figures

1.1	An instance of the envisioned wireless network architecture. . . .	4
1.2	60 GHz unlicensed spectrum allocations for different regions. . . .	6
2.1	Multiple obstacles scenario.	30
2.2	Effect of directivity on diffraction calculations.	32
2.3	WPAN simulation scenarios.	35
2.4	AP to WT LOS link loss profile.	37
2.5	Connectivity consistency from the AP to different WTs in the single hop baseline and multihop relay MAC (R denotes the primary relay: WT8 in the living room, and WT9 in the office scenario).	38
2.6	An example MAC message sequence over a superframe.	42
2.7	Aggregate network throughput for an eight WT network as a function of the number of blocked WTs connected via relays.	50
2.8	Protocol performance: living room scenario.	52
2.9	Office scenario with two relay WTs.	54
2.10	Protocol performance: office scenario.	55
2.11	Protocol performance: office scenario with two relay WTs.	56
3.1	Steerable arrays with high directivity and compact form factors.	62
3.2	Gain pattern for a flat-top antenna and a linear array of flat-top elements.	72
3.3	Approximating a circular array of slot antennas as a uniform linear array of flat-top elements.	74
3.4	The geometry of interference with directional antennas.	75
3.5	Flat-top beamwidth.	79
3.6	Flat-top antenna.	84
3.7	Linear array.	84
3.8	Aggregate network throughput versus per-flow rate.	85

3.9	Per-flow throughput empirical CDF versus node density.	85
3.10	A six node network toy example.	88
3.11	Instantaneous throughputs	89
3.12	MDMAC model: state diagram for an outgoing link.	92
3.13	An example network with 4 neighbors per node.	97
3.14	Medium utilization by node transmissions.	97
3.15	Aggregate throughput.	108
3.16	Missed transmit opportunities.	108
3.17	Throughput.	110
3.18	Total delay.	110
3.19	Delay jitter.	110
3.20	Portion of a 25 node mesh topology.	111
3.21	Instantaneous flow throughputs.	111
4.1	CSMA/CA: Backoff mechanism	124
4.2	Activity Map: Nodes A and B have a bidirectional real-time flow between them. Node C is a neighbor of nodes A and B. The activity maps of nodes A, B and C contain reservations for the real-time flow between nodes A and B. Note that the cycles need not start at the same time.	125
4.3	Nodes use CSMA to monitor the medium and mark their current Carrier Sense History Table according to the sensed medium busy and idle times.	127
4.4	At the beginning of each cycle, nodes derive their Carrier Sense Table as a summary of their Carrier Sense History Tables. The Carrier Sense Table helps to identify periodic real-time flows.	128
4.5	Establishing a new flow: The sender uses its activity map to select a suitable time window when it can establish a new real-time flow. The node transmits an R-RTS message at the chosen slot and receives an R-CTS message in response, thus setting up the flow.	131
4.6	Variable time-intervals between two consecutive backoff counter decrements at an IEEE 802.11 node.	140
4.7	The total number of voice calls supported as a function of the incoming calls for a two node network, with nodes 150m apart.	145
4.8	Line topology: the nodes are 150m apart. The edge nodes 1 and 5 act as gateway nodes, with all the calls being routed through them.	151
4.9	Grid topology: 25 nodes placed 150m apart in a square grid. The nodes colored black act as the gateway nodes for the network.	152
4.10	The effect of TXOP on the total number of voice calls supported over a five node line topology.	157

4.11 Sticky CSMA/CA call carrying capacity for the line topology. All figures show 95% confidence intervals for the average number of supported calls.	158
4.12 Instantaneous throughput of a FTP flow in a representative simulation for Sticky CSMA/CA.	161
4.13 Sticky CSMA/CA call carrying capacity for the regular grid topology. All figures show 95% confidence intervals for the average number of supported calls.	163
4.14 Sticky CSMA/CA call carrying capacity for the irregular grid topology. All figures show 95% confidence intervals for the average number of supported calls.	167
4.15 Line topology: QoS performance for voice calls.	169
4.16 Regular grid topology: QoS performance for voice calls.	169
4.17 Irregular grid topology: QoS performance for voice calls.	169
4.18 Nine node line topology: the edge nodes 1 and 9 act as gateway nodes.	170
4.19 Voice call support for nine node single line topology network. . . .	170
4.20 Voice call support for regular grid topology network with silence suppression enabled.	173
5.1 Two-scale model for resource allocation over large multihop wireless networks.	181
5.2 Linear topology.	190
5.3 Throughput profiles for mixed biasing over 50 node linear (ring) topology.	192
5.4 Regular grid topology.	194
5.5 Two-dimensional torus topology.	196
5.6 Average per-flow throughput for mixed bias with $b_1 = 5$, $b_2 = 1$, and $\alpha = 0.2$ from analysis.	200
5.7 Average per-flow throughput for mixed bias with $b_1 = 5$, $b_2 = 1$, and $\alpha = 0.2$ from simulations.	200
5.8 Scatterplots of flow throughputs for a 144-node network for mixed bias allocations, proportional fairness and severe bias ($b = 5$).	201
5.9 Scatterplots of flow throughputs for a 144-node network for mixed bias allocation policies ($b_1 = 5, b_2 = 0$).	201
5.10 Average per-flow throughput for mixed bias with $b_1 = 5, b_2 = 0$, and $\alpha = 0.2, 0.6$, and 0.8 from simulations.	205

List of Tables

2.1	Indoor setting parameters	34
2.2	Multihop relay directional MAC protocol parameters	47
3.1	MDMAC simulation parameters	106
3.2	MAC Fairness Index (MFI)	108
4.1	Overhead in terms of time, with average data rate R	142
4.2	Sticky CSMA/CA simulation parameters.	153

Chapter 1

Introduction

Over the past decade, much of the world has witnessed a wireless revolution. In their various forms, wireless networks enable untethered access to communication services. While much of the initial growth has been driven by the popularity of cellular telephony, IEEE 802.11 standards-based wireless data networks such as Wireless Fidelity (WiFi) have become very common in college campuses, enterprises, airports, cafeterias and homes as the technology of choice for Internet access. However, both cellular and WiFi networks essentially offer single hop connectivity to the wired backbone. More recently, wireless mesh networks have been proposed as a cost-effective alternative to deploying a wired infrastructure to extend Internet services to new areas. In a wireless mesh network, wireless clients connect to a wired Internet backbone via multiple hops through a backhaul mesh of wireless routers.

While wireless data rates have steadily increased through the past decade, the network usage patterns have undergone a drastic shift as well. On one hand, the number of wireless users has exploded given the decreasing costs, the higher network availability, the emergence of versatile portable computing devices such as smart-phones and laptops, and the appeal of anywhere-anytime wireless connectivity. On the other hand, with the rising popularity of applications such as high definition audio/video streaming and file sharing, the Internet data traffic is increasingly rich in bandwidth-intensive multimedia content. Consequently, current wireless networks struggle to keep up with the growing bandwidth demand, leading to an imminent need for effective solutions to this problem.

This dissertation investigates a few promising approaches to meet the increasing demand for wireless bandwidth. The key research themes that we explore are to address fundamental design bottlenecks in realizing the vision of a new generation of wireless networks with higher link bandwidths (specifically, 60 GHz networks with multiGigabit data rates), and efficient medium utilization for current and next generation wireless networks: both these themes must go hand in hand in addressing the need for higher data rates.

Before providing an overview of our contributions, a few comments on wireless system design are in order. Wireless is a broadcast medium: unlike wired communication, a transmitting node can potentially interfere with the reception of

nodes in its vicinity, causing packet loss. Therefore, for a group of wireless users operating in the same frequency band, bandwidth resources must be apportioned via appropriate resource sharing mechanisms. Medium access control (MAC) and bandwidth allocation are hence fundamental questions in wireless network design.

Several approaches can be employed to share the wireless channel among multiple users. Contention-free MAC schemes divide the wireless bandwidth (i.e., signaling dimensions) along the time, frequency, space and/or code axes, and assign dedicated channels to users. Contention-based approaches, however, require users to contend for dynamic channel access. A number of factors come into play in the design of a MAC protocol; of these, the radio propagation characteristics, the transceiver capabilities, and the network operating environment are among the most crucial. For example, a link budget analysis might reveal that the use of directional antennas is necessary to attain the desired link range and data rate combination with the allowed transmit power: directional networking comes with its own baggage of protocol design constraints. The expected driving applications and their performance requirements, as well as traffic patterns must also be given due consideration. The preceding constraints specify the design space for MAC protocols.

The overall network performance, quantified by throughput capacity, network utilization, or quality of service metrics (such as average packet-loss, delay and

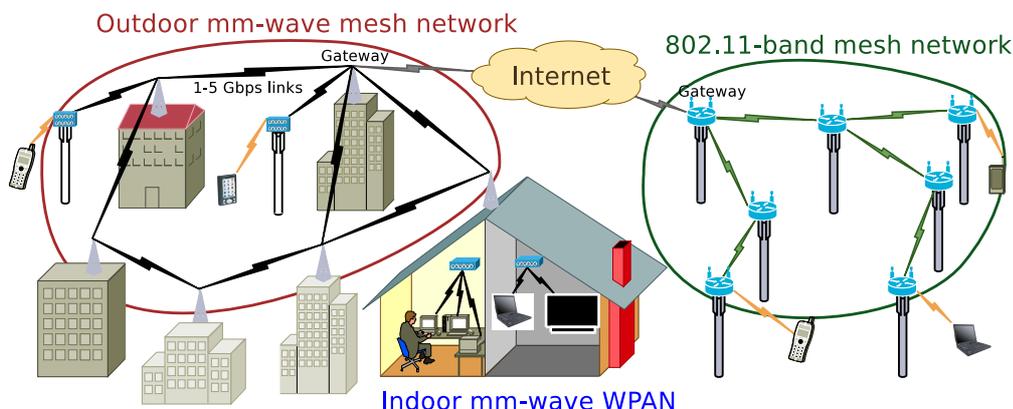


Figure 1.1: An instance of the envisioned wireless network architecture.

delay jitter), depends on how the network resources (e.g., transport capacity [1]) are partitioned among network flows. Cross-layer protocol design for efficient resource allocation among competing network connections is critical to achieve network performance goals. Indeed, this topic has been extensively studied over the last few years under the umbrella of network utility maximization (NUM), where a utility function captures the network design objective in terms of the desired tradeoff between flow throughput and fairness [2, 3].

Basically, efficient medium access control that *accounts* for, and to the extent possible, *exploits*, the wireless channel characteristics, the transceiver capabilities, and the driving network applications, along with flexible resource allocation among competing network connections, are key to enabling higher network capacity and resource utilization. This idea broadly defines the scope of this dissertation.

1.1 Dissertation Overview

Figure 1.1 depicts an instance of the envisioned network architecture comprising 60 GHz outdoor mesh networks, WiFi mesh networks, and 60 GHz indoor wireless personal area networks (WPAN). The 60 GHz outdoor mesh networks and the WiFi mesh networks operate in conjunction to provide broadband Internet access, with the lower data rate WiFi mesh networks optimized for voice over IP (VoIP) and data services. 60 GHz WPANs offer a wire-clutter-free high-definition multimedia experience at homes. Before discussing our contributions, we present a quick introduction to the 60 GHz spectrum.

The 60 GHz band, which is a part of the broader “millimeter (mm) wave” band, has been allocated worldwide for unlicensed short range wireless communications. Figure 1.2 shows the spectrum allocations in different regions. While propagation loss due to Oxygen absorption renders this band unsuitable for long range communication (e.g., of the order of kilometers) with reasonable transmit powers, this abundant unlicensed spectrum can potentially enable multiGigabit wireless communication over shorter range links. Indeed, mm wave spectrum has been used in military communications and radar systems for decades, with RFICs based on expensive packaging techniques and compound semiconductor processes. However, the use of this band is now becoming commercially attractive

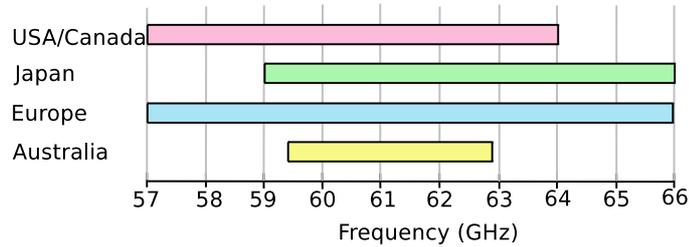


Figure 1.2: 60 GHz unlicensed spectrum allocations for different regions.

given the recent successes in building mm wave transceivers in inexpensive silicon processes, using low-cost packaging techniques [4, 5]. This has led to a significant recent interest in the use of this band to enable numerous high bandwidth indoor applications, such as high definition video streaming for wireless display/HDTV, wireless gigabit Ethernet, wireless data bus, etc.

In the first part of this dissertation, we identify and address some of the fundamental design bottlenecks in realizing the vision of 60 GHz wireless networks. In the context of 60 GHz indoor WPANs, we propose a network architecture [6] in which every link is constrained to be directional for improved power efficiency. We develop a diffraction based radio propagation model and show that obstacles such as moving human beings and furniture can result in frequent link blockages in an indoor setting. To address this problem, we propose the use of multihop relay as a practical, power efficient solution to maintain consistent network connectivity. We also define a MAC protocol that accounts for directional transmission/reception along with procedures for network discovery and recovery from link blockages.

While there is significant interest in indoor wireless personal and local area network applications of 60 GHz networking, we envision multiGigabit *outdoor* 60 GHz mesh networks with link ranges on the order of 100 meters [7]. Such networks, possibly in conjunction with longer range wireless or fiber links, can provide a quickly deployable broadband infrastructure at locations with limited wired or optical connectivity. The link spans on the order of 100m require extremely high directionality for reliable operation. Our probabilistic analysis of the spatial interference incurred due to uncoordinated transmissions shows that given the highly directional links and the high atmospheric attenuation, interference is drastically reduced, so that *pseudo-wired* link abstractions form an excellent basis for protocol design. Guided by this key insight, we propose a distributed, directional MAC protocol that employs memory and learning for implicit transmission coordination while achieving high medium utilization [8].

In the second part of this dissertation, we rethink medium access control for mesh networks operating in the 2.4 and 5.8 GHz WiFi bands. Specifically, we question the efficacy of per-packet contention design in 802.11 WiFi mesh networks for supporting a mix of real-time applications and delay-insensitive data. We propose Sticky CSMA/CA, a MAC framework that achieves time division multiplexing-like performance for real-time flows such as VoIP, without requiring network-wide synchronization [9]. We show that the implicit coordination achieved via each

node's memory of past carrier-sensed transmit activity and simple transmission rules (defined by the Sticky CSMA/CA protocol) result in significant performance gains for VoIP and data support relative to 802.11, and its QoS extension 802.11e, WLAN protocols.

Finally, we present a flexible framework for sharing transport capacity in multihop wireless networks so as to achieve a desired tradeoff between throughputs achieved by flows and their resource usage (i.e., flow throughput profiles), while ensuring efficient network utilization [10]. We show that our resource sharing mechanisms can be implemented in the form of decentralized protocols.

1.2 Contributions and Impact

In this section, we list our contributions and discuss the impact of our research in 60 GHz indoor wireless personal area networks (WPAN), 60 GHz outdoor mesh networks, real-time QoS support over WiFi mesh networks, and resource allocation in multihop wireless networks.

60 GHz Wireless Personal Area Networking [6, 11]

- We present a cross-layer modeling and design approach for multiGigabit indoor wireless personal area networks (WPANs). Our approach accounts

for the following two characteristics that sharply distinguish mm wave networking from that at lower carrier frequencies. First, mm wave links are inherently directional: directivity is required to overcome the higher path loss at smaller wavelengths, and it is feasible with compact, low-cost circuit board antenna arrays. Second, indoor mm wave links are highly susceptible to blockage because of the limited ability to diffract around obstacles such as the human body and furniture.

- We develop a diffraction-based radio propagation model to determine network link connectivity as a function of the locations of stationary and moving obstacles for a given indoor setting. This model quantifies the fragility of directional mm wave links in indoor settings with moving humans and stationary obstacles.
- For a centralized WPAN controlled by an access point, we show that multi-hop communication, with the introduction of a small number of relay nodes, is effective in maintaining network connectivity in scenarios where single-hop communication would suffer unacceptable outages.
- We define a multihop MAC protocol that accounts for the fact that every link in the WPAN is highly directional. We show via packet level simulations that our protocol maintains high network utilization with low overhead.

- Our initial work [11] is among the first published works on MAC design for mm wave WPANs and demonstrates the efficacy of multihop relay to address obstacle blockage for reliable WPAN operation. Our diffraction-based radio propagation model is a useful tool for understanding indoor mm wave radio propagation for specific indoor settings.

60 GHz Outdoor Mesh Networking [7, 8]

- We investigate an architecture for multi-Gigabit outdoor mesh networks, where mesh nodes are equipped with electronically steerable antenna arrays, with both the transmitter and receiver synthesizing narrow beams that compensate for the higher path loss at mm wave frequencies.
- Our probabilistic analysis of interference incurred due to uncoordinated transmissions shows that directionality drastically reduces spatial interference, motivating a *pseudo-wired* link abstraction as a basis for protocol design. That is, interference can essentially be ignored in MAC design, and the challenge is to schedule half-duplex transmissions in the face of the “deafness” resulting from highly directional links.
- To address the challenge of achieving transmit-receive coordination in a *distributed* fashion while achieving high medium utilization, we present Memory-guided Directional MAC (MDMAC), a distributed MAC protocol that em-

employs memory to achieve approximate time division multiplexed (TDM) schedules without explicit coordination or resource allocation.

- We demonstrate the efficacy of the MDMAC protocol via packet level simulations, while a Markov chain fixed-point analysis is employed to provide insight into the effect of parameter choices.
- This work is the first to make a strong case for multiGigabit outdoor mesh networking in the 60 GHz band. Our work spawns many open questions for further research in realizing the vision of multiGigabit outdoor 60 GHz mesh networks.

Medium Access Control for Real-time QoS in WiFi Mesh Networks [9]

- We propose a novel approach to QoS for real-time traffic over wireless mesh networks, in which application layer characteristics are exploited or shaped in the design of medium access control.
- We present Sticky CSMA/CA, a medium access mechanism that provides TDM-like performance to real-time flows without requiring explicit synchronization. Specifically, we exploit the natural periodicity of VoIP flows to obtain implicit synchronization and multiplexing gains.

- We demonstrate large gains over IEEE 802.11 and its QoS extension 802.11e in terms of increased voice call carrying capacity (more than 100% in some cases) via packet level simulations.
- This work demonstrates that memory and application layer regularities can be exploited to achieve implicit coordination without explicit network-wide synchronization. This understanding serves as the basis for even stronger application of memory and learning for coordination among *deaf* nodes in 60 GHz outdoor mesh networks [7].

Resource Biasing in Multihop Wireless Networks [10, 12]

- We propose a new class of resource allocation strategies for introducing a controlled bias against resource-intensive, long connections in a multihop wireless networks in order to significantly improve the performance of light-weight, shorter connections. Specifically, mixing strongly biased allocations with fairer allocations leads to both efficient network utilization as well as a tradeoff between flow throughput and fairness superior to that obtained by the individual strategies in the mixture.
- We present a simple analytical model that offers insight into the impact of a particular resource allocation strategy on network performance, in a manner that captures the effect of finite network size and spatial traffic

patterns. Our simulation evaluation serves to verify the analytical design prescriptions.

- By invoking a connection between our resource biasing strategies and the network utility maximization framework for resource allocation in wireless networks [2,3], we point to protocol design options to implement our resource allocation strategies.
- Our two-scale model represents a simple but powerful approach to design and performance analysis of large scale wireless networks. It facilitates a quick exploration of the rich design space of mixed bias strategies in trading off the performance seen by short and long connections.

This dissertation presents some novel ideas that have broad applications in wireless network protocol design. For instance, the use of memory and learning in MAC that we present in the context of mm wave outdoor mesh MAC and Sticky CSMA/CA is a promising approach to achieve distributed coordination in a large wireless network with little overhead. The inspiration for this idea comes from the biological phenomenon of *stigmergy*, which, broadly speaking, is a mechanism for implicit communication where each entity alters its behavior in reaction to perceived changes in its local environment caused by other individuals, obviating the need for explicit communication. This interaction can facilitate emergent

coordination among networked entities to lead to a desired global behavior without the need for explicit communication, centralized control or planning. Example outcomes of stigmergic coordination in nature are plenty, including complex structures built by social insects such as termite mounds and ant hills, and ant trails to food sources that evolve towards shortest paths. This idea is known to have profound applications in a number of areas such as multi-agent systems, robotics, optimization, data mining and communication networks, to name a few.

Our work in MAC design for 60 GHz outdoor and indoor networks is among the few early works to consider the design of fully directional MAC protocols, without an optional omni-directional mode generally exploited for control in many prior directional networking proposals in the lower frequency WiFi (2.4 and 5.8 GHz) band. Further, these works serve to bring the networking community's attention to the promise of 60 GHz networking and several research challenges of this exciting new frontier in high-speed wireless networking.

The approach to MAC protocol design presented in this dissertation illustrates the importance of cross-layer considerations in wireless network protocol design. For instance, Sticky CSMA/CA achieves implicit coordination and TDM-like resource sharing by employing periodic transmission schedules (that exploit natural or imposed application layer regularities) in conjunction with enabling nodes to learn, and adapt to transmit patterns in their neighborhood by remembering the

physical carrier-sensed transmission activity. The design of 60 GHz indoor WPAN MAC protocol is driven by two key physical characteristics of mm wave links in indoor settings: directionality, and link blockage from obstacles, which we rigorously study via the diffraction-based indoor mm wave radio propagation model before embarking on the design of our directional multihop relay MAC protocol. Similarly, 60 GHz outdoor mesh MAC protocol is shaped by two crucial design implications of using highly directional mm wave links: (1) deafness, which makes MAC based on carrier sensing infeasible, and transmitter-receiver coordination extremely challenging; and (2) drastic reduction of spatial interference because of high directionality and Oxygen absorption, which motivates pseudo-wired link abstraction as a basis for protocol design.

1.3 Dissertation Outline

The remainder of the dissertation is organized as follows. Chapter 2 presents a cross-layer approach to MAC design for 60 GHz WPANs. In Chapter 3, we motivate 60 GHz outdoor mesh networks and show that for a rooftop type deployment scenario, a pseudo-wired model of wireless links suffices as a first order approximation for MAC design. We then present a MAC protocol for outdoor mm wave mesh networks. We describe Sticky CSMA/CA, a MAC protocol for WiFi mesh net-

works in Chapter 4. In Chapter 5, we present a framework for resource allocation over multihop wireless networks and propose new resource biasing strategies that mix strongly biased allocations with fairer allocations to shape throughput profiles in more flexible ways than that offered by the prior state-of-the-art. Chapter 6 concludes this dissertation with a summary of contributions and open issues.

Chapter 2

60 GHz Wireless Personal Area Networks

2.1 Introduction

In this chapter, we investigate indoor wireless personal area networks (WPANs) utilizing the 60 GHz “millimeter (mm) wave” band. The recent trends in low-cost mm wave radio frequency (RF) integrated circuits (ICs) designs and the ever growing demands for high bandwidth fueled by surge of multimedia applications have spurred intense efforts in both research and standardization for indoor mm wave communication [4, 14–21]. However, the eventual success of these efforts depends on system designs that account for the fundamental differences between mm wave communication and existing wireless networks at lower carrier frequencies (e.g., from 900 MHz to 5 GHz). In particular, the goal of this chapter is to present a

cross-layer modeling and design approach that addresses two unique features of mm wave links: their vulnerability to blockage and their inherent directivity.

Blockage: Electromagnetic waves have very limited ability to diffract around obstacles whose size is significantly larger than the wavelength. Since the wavelength at 60 GHz is so small (5 mm), links are effectively blocked by obstacles such as humans and furniture. For example, blockage by a human can penalize the link budget by 20-30 dB. Human movement in a room, therefore, can cause intermittent blockage of mm wave WPAN links, resulting in a time-varying network topology. Maintaining seamless network connectivity and providing the QoS needed for applications such as HDTV in such a setting is a challenge for network protocol design which is fundamentally different from that at lower carrier frequencies. Meeting this challenge requires models that take into account diffraction (or the lack thereof), both for obtaining design insight and for performance evaluation.

Directivity: Millimeter wave links are inherently directional. The free space propagation loss between isotropic antennas scales as λ^2 , where λ is the carrier wavelength, so that 60 GHz is 21.6 dB *worse* than 5 GHz for omnidirectional communication. On the other hand, for a fixed antenna aperture area (which depends on the transceiver form factor), directivity scales as $1/\lambda^2$, giving a gain of $1/\lambda^4$ when we account for both transmit and receive antennas. This corresponds

to an overall scaling of $1/\lambda^2$, so that a 60 GHz link with directional transmission and reception can now become 21.6 dB *better* than a 5 GHz link, assuming nodes of similar form factor. In particular, steerable antenna arrays implemented as patterns of metal on circuit board can be employed to synthesize highly directional beams for 60 GHz nodes that are smaller than a typical WiFi access point. It is crucial to take advantage of this, given the difficulty of producing a large amount of transmit power at mm wave frequencies using low-cost, low-power silicon implementations.

Our approach: We propose a cross-layer modeling framework and a multihop directional MAC architecture for robust, multiGigabit, in-room WPANs. Each node has an electronically steerable directional antenna, so that the transmitters and receivers can steer beams towards each other. The key idea is to handle blockages by going around obstacles, rather than burning through them. Blockages that result in a 20 dB power loss require a 100-fold reduction in data rate in order to maintain the same reliability, when operating in a power-limited regime. On the other hand, routing around the obstacle by replacing the blocked link by 2 links only reduces throughput by a factor of two. We consider directional, line of sight (LOS) links, with each link operating at a fixed nominal data rate (e.g., 2 Gbps) when the LOS path is available. When the LOS path between two nodes is blocked, we route around it, still using directional LOS links. Assuming that

there are enough spatially dispersed nodes, this multihop architecture provides both high power efficiency and robust connectivity in the face of stationary and moving obstacles typical of living room and office settings. In principle, it is also possible to use reflections from walls and other surfaces to steer around obstacles; however, not only do reflections result in a loss of link budget, but the efficacy of using them to avoid blockage is a sensitive function of node placement and the propagation environment. We therefore focus on whether it is possible to obtain robust network connectivity and high throughput with LOS links alone, using a small number of relays if necessary.

We consider the specific in-room scenario of several wireless terminals (WTs) controlled by a single access point (AP), with nominal operation consisting of direct, contention-free communication between the AP and each WT. When an AP discovers that its link to a WT is blocked, it sets up an alternate route based on its current topology information (the topology of the directional links is discovered during set-up, and regular opportunities for topology update are a part of the proposed MAC protocol). Due to the slow time scale of human movements (which are typically the cause of network topology changes) relative to the topology updates enabled by the protocol, our simulations show that the alternate routes computed by the AP are invariably functional. Thus, while outage rates for a given link can be quite high (as high as 60%), intelligent multihop

networking effectively removes outages in connectivity between the AP and the WTs. An important component of our work is the development of simple models for time-varying blockage in typical WPAN environments that enable performance evaluation of the preceding architecture.

Summary of contributions: Our contributions are summarized as follows:

- 1) We propose and investigate an in-room WPAN architecture that addresses both blockage and directivity, the fundamental features that distinguish mm wave networks from those at lower frequencies. In particular, the proposed directional MAC protocol is designed for a network in which every link is constrained to be directional, without fallback to an omnidirectional mode for coordination as in most prior work. The protocol includes procedures for topology discovery and updates, and recovery from LOS link outages via multihop relay to the blocked nodes.
- 2) We analyze the effect of obstacles on the received signal strength via a site-specific mm wave propagation model based on the Fresnel-Kirchhoff diffraction theory. This provides a simple model to track the time evolution of link losses using deterministic computations, and hence the network connectivity for a given set of stationary and mobile obstacles whose geometry models a human.
- 3) We evaluate the performance of our multihop relay directional MAC protocol via analysis and extensive packet level simulations. The simulations verify the

efficacy of multihop relay in maintaining consistently high throughput with low control overhead, despite frequent LOS link outages due to time-varying blockage.

2.2 Related Work

To the best of our knowledge, there is no prior work on the design of mm wave WPANs with exclusively directional links in the literature, except for a conference publication with our own preliminary results [11]. There are many ongoing industry-led efforts aimed towards the standardization of 60 GHz WPAN network interfaces; for example, the IEEE 802.15 WPAN Millimeter Wave Alternative PHY Task Group 3c [19], the WirelessHD Consortium [20], and ECMA International [21]. The ECMA-387 specification for the 60 GHz WPAN PHY/MAC released in December 2008 also includes a relay mechanism to counter link blockage thereby affirming the timeliness of this work and the need for further research to better understand the efficacy of such mechanisms.

The use of directional antennas has been extensively studied for wireless networks operating in the lower frequency bands such as cellular and broadband networks, and over the last decade, on WiFi-based multihop networks [22–28]. In this context, the performance benefits of directional communication such as improved spatial reuse and extended directional communication range have an associated

cost, because a number of problems arise due to, or are aggravated by, directional communication, such as the hidden terminal problem, deafness [27], and the problem of neighbor discovery. The typical solution in the literature is to employ a combination of directional and omnidirectional communication for critical control message exchanges (e.g., the four way handshake in IEEE 802.11) for medium access control. Such a dual-mode operation is not appropriate for the mm-wave WPANs that we envision, for which directionality is required at both the transmitter and the receiver simply to achieve reliable high data rate communication. For such networks, network protocols based on a directional mode alone (with the ability to choose the direction, either via sectorization or beamsteering) need to be developed. References [28, 29] propose fully-directional MAC protocols for multi-hop wireless networks: [28] proposes a directional slotted ALOHA protocol that exploits the adaptive beamforming capabilities of smart antenna arrays, whereas [29] presents a polling-based decentralized MAC protocol. However, none of the papers in the literature model or address the problem of frequent link outages due to blockage, which fundamentally alters the design tradeoffs for mm wave networks relative to those at lower carrier frequencies. In particular, far more agility needs to be designed into the network protocols to handle the time-varying network topology with low overhead, while satisfying the stringent QoS requirements for the bandwidth-hungry applications driving the development of mm wave WPANs.

Network design efforts such as ours are motivated by recent advances in mm wave circuit design [4, 14, 16], including multiGigabit electronically steerable directional links [5], that indicate that low-cost commercially feasible realizations are within reach. In terms of the channel model we use, while our diffraction model is based on fundamental physics, we are motivated by the extensive body of knowledge on mm wave propagation measurement and modeling. Measurement campaigns in indoor environments include [30–39]. For typical indoor environments with omnidirectional antennas, specular reflections from surfaces are dominant contributors to the received signal power as compared with diffraction or scattering [34, 40–42]. Since the path that is the strongest in such a setting is the LOS path (if it is not blocked), this motivates restriction to LOS for maximizing power efficiency. The reduction of multipath for directional mm wave links [33, 42, 43] means that link budget calculations for a simple additive white Gaussian noise channel model are reasonably accurate for a directional LOS link. The susceptibility of mm wave links to blockage due to their weak diffraction characteristics is well known [33, 44], and the effect of human movement is investigated in [45, 46], but their impact on the network performance has not been studied previously. Many deterministic and statistical mm wave propagation models have been proposed based on channel measurement studies [35, 40, 47], but many of these focus on omnidirectional transmission (and possibly directional reception). Reference [48]

provides line of sight (LOS)/ non-line of sight (NLOS) channel models developed by the IEEE 802.15 TG3c group for indoor WPAN environments. These statistical channel models do not account for the effect of moving obstacles such as humans on the network connectivity over time (e.g., deep fades for seconds), and are basically meant to be used for a comparison between different physical layer designs.

Outline: We describe our physical layer model in Section 2.3, including a model for blockage by both stationary and moving obstacles. Simulation results with this model are used to motivate the need for multihop communication to provide robust network connectivity. Section 2.4 presents a multihop directional MAC for achieving such robustness, and estimates achievable rates accounting for the overhead. In Section 2.5, packet-level simulations of MAC performance, taking into account the blockage model developed in Section 2.3, are used to demonstrate the efficacy of our multihop architecture. Section 2.6 contains our conclusions, including a brief discussion of important areas for future research.

2.3 Physical Layer Model

We first describe an example link budget for a LOS 60 GHz link to give a feel for the required transmit signal power levels for the feasibility of WPANs with di-

rectional LOS links. However, we then abstract away from detailed design choices in the physical layer to focus on the key bottleneck for mm wave communication: blockage by obstacles. We describe in detail the calculations needed to compute diffraction-based path loss for a given obstacle configuration, and use the model to obtain the time-varying network connectivity. We show that, while any given link can frequently be in outage due to blockage, multihop communication can indeed provide robust connectivity. This motivates the multihop MAC in the next section.

Example Link Budget: We present simple calculations that indicate the feasibility of a plug-and-play WPAN. The directivity of an antenna is the ratio of the maximum power density (watts/m²) to its average value over a sphere. The directivity of an antenna can be approximated as [49]:

$$D = \frac{40000}{\theta_{HP}^o \phi_{HP}^o}$$

where θ_{HP}^o and ϕ_{HP}^o are the horizontal and vertical beamwidths, respectively, of the antenna. For a WPAN application, we might design an antenna element to have a horizontal beamwidth of 120° and a vertical beamwidth of 60°, which allows a rough placement of nodes in order to ensure LOS to one or two neighbors. The directivity for such an element, which can be realized as a pattern of metal on circuit board, is 5.55 (or 7.4 dBi). If we put four such elements to form a steerable antenna array, we can get a directivity of 22 (or 13.4 dBi). Now, assuming an an-

tenna directivity of 13.4 dBi at each end, we do a link budget for a QPSK system operating at 2 Gbps [50]. For a receiver noise figure of 6 dB, bit error rate of 10^{-9} , excess bandwidth of 50%, and assuming free space propagation, we obtain that the required transmit power for a nominal range of 10 meters is about 8 mW, including a 10 dB link margin. When split among four antenna elements, this transmit power corresponds to 2 mW of power per antenna element. RF front ends for obtaining these power levels are realizable with silicon semiconductor processes, thereby indicating the feasibility of low-cost, high-volume production of the kinds of WPAN nodes on which our architecture is based.

Adaptive Beamforming Antennas: We assume that all the WPAN nodes are equipped with electronically steerable antenna arrays that can be used to provide directivity on both transmit and receive. A receiving antenna array uses a training sequence in the PHY preamble of a packet to adjust the required array weights in order to adaptively beamform towards the direction of the signal of interest [51], using standard adaptive algorithms [52], such as least mean squares (LMS) or recursive least squares (RLS), or variants thereof. For a relatively small number of elements (e.g., 4-10), such algorithms converge quickly, so that packet-by-packet beamsteering with, say, 50 symbols training overhead, should be feasible. Thus, we assume that an idle node can quickly steer its receive antenna array towards an incoming signal: this amounts to “omnidirectional” sensing (restricted only

by the field of view of the antennas in the array), but directional reception. The beamforming weights learnt when receiving from a given node can then be used to transmit back to that node, using channel reciprocity, so that transmission is directional as well. Alternatively, the transmitter or receiver could choose from among a fixed number of sectors, thus discretizing the beamforming function.

Diffraction due to obstacles: We use a simple geometric model to estimate the diffraction loss along the LOS path between two nodes, taking into account the node placements, the locations and dimensions of obstacles, and the room dimensions. We neglect the contribution from the reflected signals to the received signal power; narrow beam directional antennas along the LOS direction substantially reduce the contribution of reflected multipath components [33, 42, 43, 46].

We make the following simplifying assumptions in modeling obstacles:

- 1) The attenuation due to an obstacle in the LOS path is so high that the energy of the signal propagating through the obstacle is negligible. In other words, we only consider the obstacles that can cause a significant attenuation to a signal propagating through them. For mm waves, most of the common obstructions in indoor environments, such as human beings, thick walls and furniture, fall in this category. Thus, the link gain is only due to diffraction around the obstacle.
- 2) The human body is approximated as a perfect conducting cylinder, whose pro-

jection on the plane perpendicular to the line of sight is considered for diffraction calculations. Other obstacles are approximated in a similar manner.

Diffraction of electromagnetic waves [53, 54] can be intuitively explained in terms of a fundamental principle from physical optics: the Huygens' principle which states that *“each point on a primary wavefront can be considered to be a new source of a secondary spherical wave and a secondary wavefront can be constructed as the envelope of these secondary spherical waves”* [54]. An obstacle blocks a subset of these secondary waves, and the field at a point of interest beyond the obstacle can be obtained by summing up the contributions from the remaining secondary waves. The mathematics of these computations is often referred to as Kirchhoff theory [53]. Diffraction theory has been widely used to study terrain diffraction in the context of wireless cellular systems primarily using the knife-edge diffraction model [55–59].

Fig. 2.1 shows the propagation path from a source Tx to a receiver Rx obstructed by n obstacles, modeled as perfect conducting strips. To begin with, assume the transmitter and the receiver to be isotropic point sources (we specify later how to incorporate directivity). We now calculate the diffracted electric field at the receiver by successively applying the Huygens' principle at each obstacle along the LOS from the transmitter to the receiver. This is done by summing up contributions from “non-blocked” secondary sources at one obstacle, at the plane

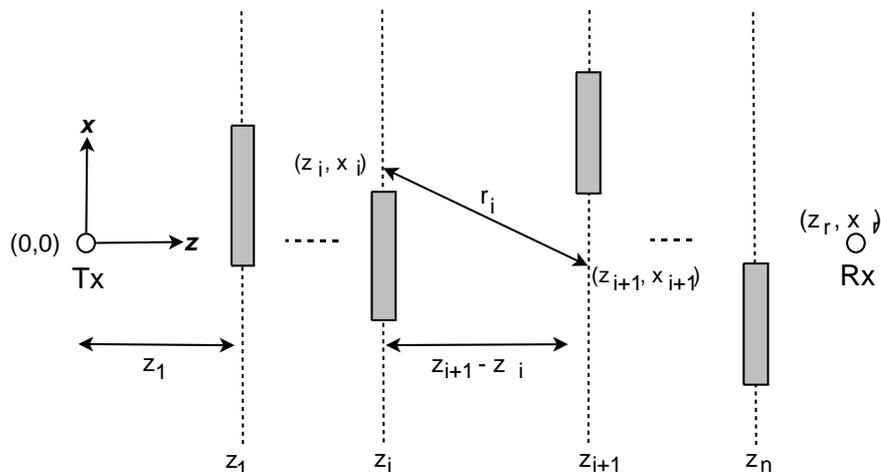


Figure 2.1: Multiple obstacles scenario.

of the next obstacle. These iterative computations are initialized by specifying the electric field $E_1(x_1)$ at a point (z_1, x_1) on the plane containing obstacle 1:

$$E_1(x_1) = \frac{E_c}{r_1} e^{-j\beta r_1} \quad (2.1)$$

where E_c is a constant, $r_1 = \sqrt{z_1^2 + x_1^2}$ and $\beta = \frac{2\pi}{\lambda}$ is the phase constant for wavelength λ (the time variation $e^{-j\omega t}$ is suppressed).

We now compute the gain from the i th plane to the $(i+1)$ th plane of interest, where $i = 1, \dots, n-1$. The same computation applied for $i = n$ gives the field at the receiver, whose location corresponds to the $(n+1)$ th plane. Let us now specify a typical step (say step i) in these calculations. The electric field $E_{i+1}(x)$ at a point (z_{i+1}, x_{i+1}) in the $(i+1)$ th plane is calculated by superimposing the contributions from the Huygens' point sources at points (z', x') in the i th window,

as follows:

$$E_{i+1}(x) = \int_{-\infty}^{\infty} h_i(x')E_i(x')g_i(x-x')dx' \quad (2.2)$$

where function $h_i(x') = e^{-j\beta(z-z')} \sqrt{\frac{jz'}{\lambda z(z-z')}} I_i(x')$, $g_i(x') = e^{-j\beta \frac{x'^2}{2(z-z')}}$ and $I_i(x')$ is an indicator function:

$$I_i(x') = \begin{cases} 1 & x' \in \{\text{Obstacle } i\}' \\ 0 & x' \in \{\text{Obstacle } i\}, \end{cases}$$

While we state the preceding results without proof, we refer to [55] for detailed derivations of similar formulas in the context of terrain diffraction where obstacles are treated as knife-edges. Note that the y coordinate perpendicular to the plane of the chapter has already been integrated out under the assumption that the obstacle heights extend beyond the first few critical Fresnel zones along the y dimension.

Equation (2.2) is the convolution of functions $f_i(x) = h_i(x)E_i(x)$ and $g_i(x)$. Successive convolutions as we go from obstacle to obstacle can be efficiently computed using the Fourier transform, since they correspond to multiplications in the frequency domain. In order to obtain accurate field estimates using the FFT and IFFT methods, it is important to choose the spatial sampling intervals such that the aliasing errors are minimized. The size of the computation window should be chosen such that the secondary wave sources outside the computation window do not have a significant effect on the resulting electric field at the receiver. For our

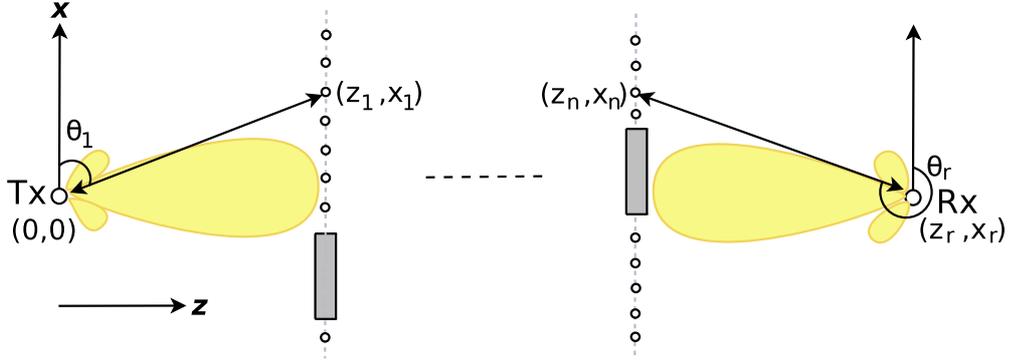


Figure 2.2: Effect of directivity on diffraction calculations.

example indoor scenarios, we use a computation window of size 2m (400λ) with a 4096 point FFT (sampling interval $< 0.1\lambda$). Also, the minimum gap ($z_{i+1} - z_i$) between two computation windows i and $i+1$ is assumed to be $\geq 25\text{cm}$ for better accuracy of the field predictions that require the distance approximations in the Fresnel-Kirchhoff formulation to hold.

We now show how to modify these computations to account for directivity. Transmit directivity means that the initialization (2.1) must be modified to account for the transmit antenna array radiation pattern. Referring to Fig. 2.2, the electric field at a point (z_1, x_1) is given by

$$E_1(x_1) = \frac{E_c}{r_1} e^{-j\beta r_1} A_{tx}(\theta_1(z_1, x_1)), \quad (2.3)$$

where $r_1 = \sqrt{z_1^2 + x_1^2}$, $\frac{E_c}{r_1} e^{-j\beta r_1}$ is the field at (z_1, x_1) due to an isotropic transmit antenna and $A_{tx}(\theta_1(z_1, x_1))$ is the transmit antenna array pattern for an azimuthal angle $\theta_1(z_1, x_1)$ and vertical angle $\phi = 90^\circ$. The antenna array pattern for different

directions can be calculated from the number and placement of antenna array elements, the radiation pattern of elements, and the input array weights [53]. Receive directivity requires the modification of the last step in the computations, as follows: Specifically, the superposition of the Huygens' point sources in the plane of the n obstacle at the receiver must account for the receive antenna pattern, as follows:

$$E_{n+1}(x) = \int_{-\infty}^{\infty} h_n(x')E_n(x')A_{rx}(\theta_r(z', x'))g_n(x - x')dx' \quad (2.4)$$

Here $A_{rx}(\theta_r(z', x'))$ is the receive antenna array field pattern for an azimuthal angle $\theta_r(z', x')$ and vertical angle $\phi = 90^\circ$. The preceding method accounts for directivity more accurately than the simpler technique [56] that we used in our preliminary results reported in [11], where we calculate the diffraction loss in dB for an isotropic transmission and reception, and add the antenna gains later.

Finally, we express the diffraction loss relative to the free space propagation loss without obstacles. Letting $E = E_{n+1}(x_r)$ denote the field obtained at the receiver, we define the diffraction coefficient $D = \frac{E}{E_{fs}}$, where E_{fs} is the electric field at the same point assuming unobstructed free space propagation. The latter is given by $E_{fs} = \frac{E_c}{d}e^{-j\beta d}A_tA_r$, where E_c is a constant, d is the distance between the transmitter and the receiver, $\beta = \frac{2\pi}{\lambda}$ is the phase constant for wavelength λ , and $A_t = \max_{\theta}A_{tx}(\theta)$ and $A_r = \max_{\theta}A_{rx}(\theta)$ are the transmit and receive antenna directivities, respectively, assuming that the transmit and receive antenna arrays

Parameter	Value
Human height range	(1.5m - 2.1m)
RWP model: velocities (min,max), pause time	(0m/s,1m/s), 10s
Fixed obstacle height range	(1m - 1.4m)
WT location height range	(0.5m - 1.5m)
AP location heights home/office	2m/2.5m
Sampling time interval	100ms

Table 2.1: Indoor setting parameters

are optimally oriented towards each other. The (relative) diffraction loss in dB is given by $L_{dB} = -10 \log_{10} |D|^2$. In order to determine the overall path loss, we simply add it to the unobstructed free space propagation loss in dB .

Indoor Radio Propagation Simulation: Based on the diffraction model described in this section, we have developed a MatLab radio propagation tool to evaluate the link losses between different nodes in a given indoor environment with human beings and other obstacles. This tool yields link losses between different network nodes as a function of time. The inputs to the tool are the parameters required to simulate a WPAN in a specified 3-dimensional indoor environment: the room dimensions; the number, position, and dimensions of the stationary obstacles such as furniture; the number of human beings; the placement of the AP; the number and positions of the WTs; and the antenna configurations. We use the Random Waypoint (RWP) model [60] for human movements in the room. We assume that all the nodes are equipped with a linear, adaptive beamforming

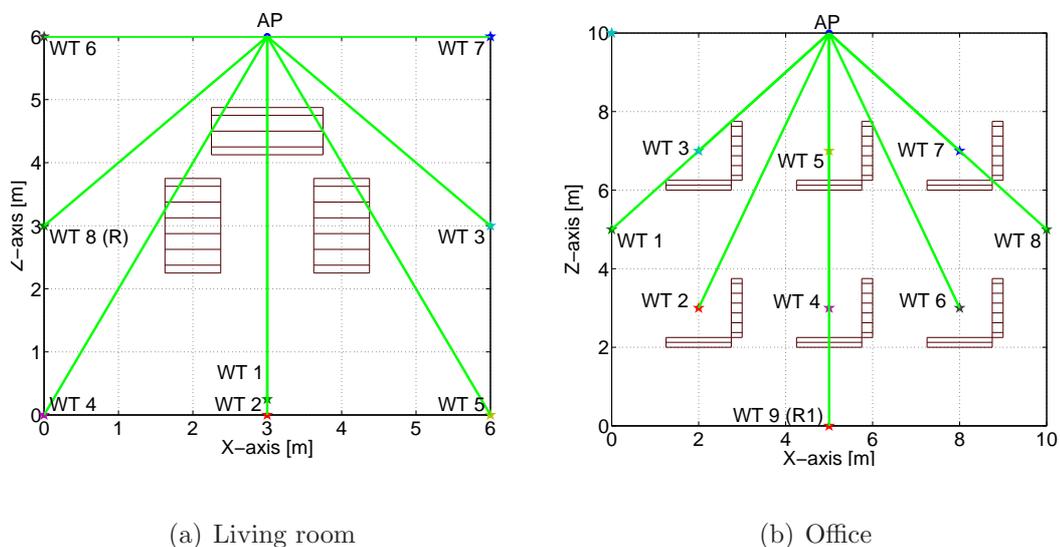


Figure 2.3: WPAN simulation scenarios.

array of 16 isotropic antenna elements spaced $\lambda/2$ apart. The default configuration parameters related to the test scenarios are listed in Table 2.1.

We now use our radio propagation tool to quantify network connectivity for some example WPAN scenarios. We consider two different indoor settings that model typical scenarios where 60 GHz WPANs are expected to be deployed: a living room and an office space (see Figs. 2.3(a) and 2.3(b)). The living room scenario has a WPAN formed by an HDTV, a surround sound system with speakers at room corners and a desktop/printer; and has eight human beings, i.e., during a gathering at home. The office space scenario has desktops and printers forming a WPAN, with fifteen human beings. The room and obstacle dimensions and the node placements have been chosen as representative of the real world scenarios

in which a large number of people can cause a high blockage probability for the individual links. Note that WT8 in the home scenario and WT9 in the office scenario are placed higher (2.5m) than the other WTs such that they have a high probability of a clear LOS connectivity to most of the WTs and the AP. Hence they can act as effective relays in case the direct LOS connectivity from the AP to a WT is blocked.

We define *connectivity consistency* as the percentage of time out of the total operation period of the network when a WT is reachable from the AP either through a direct LOS link or through a multihop path consisting of live direct links. For illustration, we employ the following link outage criterion for calculation of the expected connectivity consistency: if the diffraction loss due to obstacles exceeds 10 dB for a link, then it is considered to be in outage. This model is pessimistic because link budgets are determined based on a maximum range of operation (10 meters for in-room operation). Thus, links over shorter ranges may have enough link margin to “burn through” the obstacles. By abstracting away the dependence of connectivity on range, we obtain a worst-case network connectivity estimate that serves to stress-test the proposed multihop architecture.

We note that connectivity consistency is an indirect metric because the need for connectivity arises only if there is data to transmit at either side. However, considering the importance of an LOS link for maintaining direct connectivity

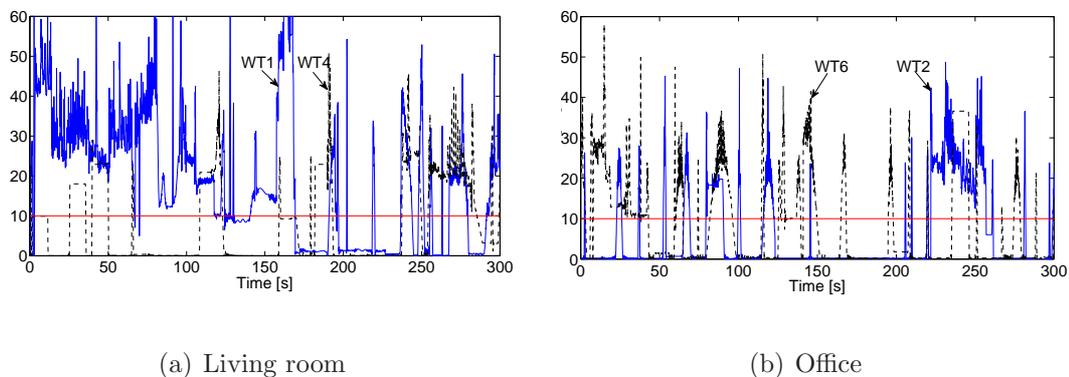


Figure 2.4: AP to WT LOS link loss profile.

(unlike for 2.4 or 5 GHz WLANs), the links can become blocked easily due to obstacles in indoor environments. Thus, this metric characterizes the actual connection state and the data transfer capacity of the network. This metric can also be interpreted as an indicator of the maximum aggregate throughput sustained by the network when all the nodes have data to send.

Figs. 2.4(a) and 2.4(b) plot the variation of the diffraction loss for some specific WT links as a function of time for the living room and office environments over a sample period of 300 seconds. We observe that there are heavy link losses because of the large number of human beings (and their random movements) and the stationary furniture obstacles in both the environments. These obstacles result in intermittent connectivity to the affected WT if the underlying MAC completely relies on the direct single hop connectivity of the AP to the WTs. These loss results demonstrate that networks with the baseline single hop MAC schemes will not

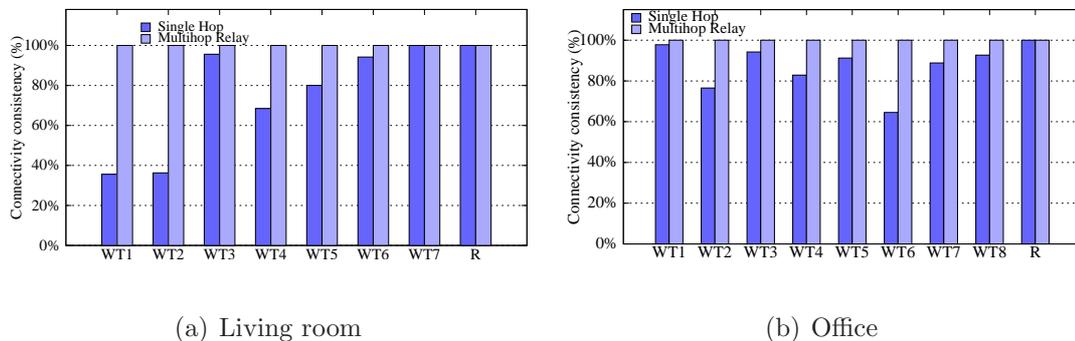


Figure 2.5: Connectivity consistency from the AP to different WTs in the single hop baseline and multihop relay MAC (R denotes the primary relay: WT8 in the living room, and WT9 in the office scenario).

be able to provide the required QoS guarantees to different WPAN applications, which are essential requirements for any practical WPAN solution.

Fig. 2.5 compares the expected connectivity consistency between single hop communication-based approach and a multihop communication scheme that exploits relays to reach the blocked WTs. We observe that, on average, the connectivity consistency for the baseline single-hop communication is significantly lower than the multihop relay scheme, which is able to maintain almost 100% network connectivity by using multihop paths via relay node(s) to connect to the blocked node(s).

Note that the high availability of alternate routes in a multihop architecture can be assured by an appropriate placement of the relay nodes (e.g., high up on the walls, or on the ceiling) accounting for the vertical beamwidths of the

WT antennas such that the relays are readily accessible when needed. On the other hand, the poor connectivity consistency of single hop communication makes it unsuitable for WPAN applications with stringent QoS requirements such as multimedia streaming. It is interesting to note that the non-zero single hop connectivity consistency data for both the living room and the office setting implies that no WT is completely blocked by the stationary obstacles. Thus, the moving obstacles (humans) play a key role in causing blockages in both example settings considered. The stationary obstacles constrain human movements and the choice of relays for different WTs because many WTs are blocked from each other.

Having shown that multihop relay helps to maintain consistent connectivity on the face of frequent link blockages due to obstacles, we next present a WPAN MAC protocol that accounts for directional communication and has built-in intelligence for topology discovery and recovery from link blockages via multihop relay.

2.4 Directional MAC Design

The key idea behind our multihop relay directional MAC framework is to utilize a mix of the conventional AP-based single hop MAC architecture for primary connectivity and resort to the multihop ad hoc mode with intermediate nodes acting as relays (though still controlled by the AP) to prevent drastic reduction of

data rates or link outage when the LOS path to a WT is obstructed. We briefly outline a mechanism for network topology discovery first and then describe the main components of our MAC protocol.

2.4.1 Discovery Algorithm

During the network initialization phase, the AP sends a Hello message and waits for the response from the WTs in each sector (the geographical region around a node is divided into equal angular sectors based on the antenna array horizontal beamwidth). The preamble preceding the Hello message payload acts as a training sequence for the WTs to beamform in the AP's direction. The WTs that successfully receive the Hello message record the antenna array weights corresponding to the direction of the AP in their network topology map and use the same weights to respond to the AP, using reciprocity. The unregistered WTs in a sector employ a Slotted Aloha contention scheme for transmitting the Hello Response message over the next m slots following the receipt of the Hello message, with the Hello Response transmission probabilities dictated by the AP (via the Hello and Hello Response ACK messages). Here, a slot duration is sufficient for the transmission of a Hello Response message and the corresponding Hello Response ACK message from the AP to confirm a successful registration. For narrow angular sectors, there is likely to be few WTs in a sector, so the problem of excess collisions is

unlikely. After completing a round of discovery with a nominal Hello Response transmit probability (e.g., 0.2 – 0.5), this procedure can be repeated with a Hello Response transmit probability of one to verify that no WT is left unregistered.

After performing the discovery procedure and having formed a network topology map (i.e., the identities of the WTs in the network and the appropriate antenna array weights required to reach them), the AP iteratively designates each WT among the registered nodes to perform the same discovery procedure. This process continues until all the WTs have finished the discovery procedure and have created their own network topology maps such that each WT is aware of the identities of the other WTs and the appropriate antenna array configurations required to reach them. Every WT sends its network discovery table to the AP after it completes its network discovery process. This information helps the AP deduce the link connectivity status of the whole network, which is useful in the determination of a relay WT when a the direct link to a node is blocked, as described later in this section.

2.4.2 Normal Mode of Operation

The AP sequentially polls all the registered WTs to check connectivity to each WT and to check whether any WT has data to transmit. Each WT must respond within a fixed interval, i.e., Poll Inter Frame Space (PIFS), with a data packet or with a *connection live* poll response message if it does not have any data to trans-

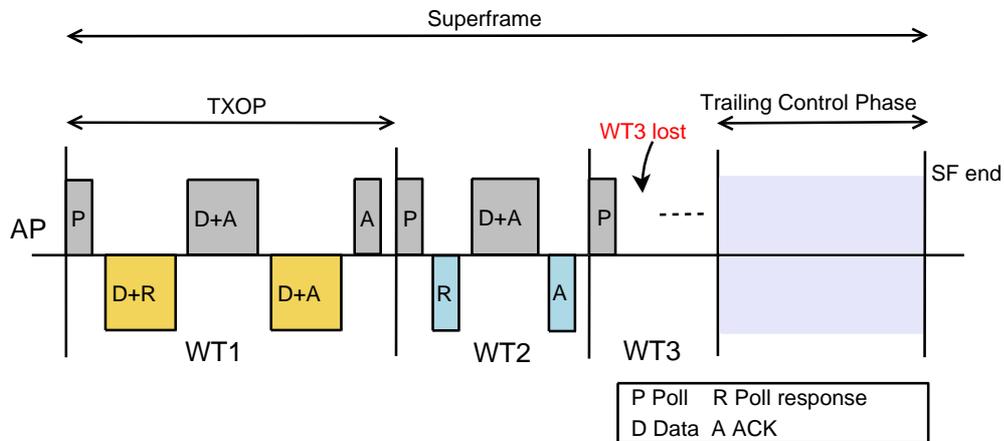


Figure 2.6: An example MAC message sequence over a superframe.

mit. This polling-based procedure is required because the conventional carrier sensing medium contention schemes do not work well with directional antennas. The polling scheme helps the AP to track the connectivity to all the registered WTs: the absence of a poll response from a WT is assumed to indicate link blockage, and triggers the search for a relay node to set up a multihop relay path to the WT. Note that with a link budget that ensures reliable link operation over the desired range, and in the absence of simultaneous transmissions that can cause interference loss, link blockage from obstacles is the dominant cause of packet loss in this setting. However, interference loss increases when simultaneous transmissions are allowed within the network to achieve high spatial reuse, or in the case of high-density WPAN deployments where multiple networks must coexist in close proximity. The use of a relay node (in addition to other mechanisms such as rate

control) can help counter packet loss even in these scenarios. However, additional procedures need to be defined to enable high spatial reuse within a WPAN and coexistence of multiple WPANs. These procedures are topics of future research, and are beyond the scope of this work.

The WT can continue to receive or send data packets until a maximum allowed time duration called the *transmission opportunity* (TXOP) duration. Thus, data transmission in both directions can be bursty because of the sequential transmission of multiple packets in response to the poll message, up to the TXOP duration. Besides providing better QoS performance for the inherently bursty multimedia streaming applications, this allows the WTs to better utilize the available LOS connectivity and also minimizes the control overhead associated with data packet transmissions. If the AP sends a data packet to a WT, the WT acknowledges the successful packet reception either by piggybacking an ACK message on the next data packet that it has for the AP or by sending a separate ACK message.

The AP's dwell time in each sector depends on the data transmission requirements of the WTs in that sector. The AP to WT data transmission mechanism follows a weighted round robin scheduling approach such that the desired level of QoS to different WTs is ensured.

2.4.3 Trailing Control Phase

The trailing control phase is utilized by the AP to allow new WTs to register and perform a network discovery procedure while the network is operational. During the trailing control phase, the AP can also verify its own topology map or designate the registered WTs to verify their network topology maps by sequentially sending Hello messages to each WT. The trailing control phase is limited to a maximum duration, which is larger than the average successful discovery phase time of a node. Because the regular network topology verification procedure of the trailing control phase occurs at a rate much faster than the dynamics of the indoor environments (human movements or change in the room setup), the AP is aware of the LOS connectivity of all the WTs and it can use the topology verification/discovery reports sent back by the WTs to choose a candidate relay node for a blocked WT.

A superframe is defined as the time taken by the AP to poll all the registered WTs in the network. The maximum superframe duration is limited by the number of WTs in the network, the TXOP duration, and the trailing control phase duration. Fig. 2.6 illustrates an example data transmission and control message sequence over a superframe.

2.4.4 Lost Node Discovery and Establishing a Relay Path

If the AP does not receive a poll response from a registered WT, it considers the WT to be *lost* and intelligently chooses a WT among the *live* WTs (with expected LOS connectivity to the lost WT as determined from the regular topology verification reports from the WTs) to act as a relay to the lost node. It commands the chosen relay WT via a Search Lost WT Request message to discover (i.e., check connectivity status with) the lost WT and report back within a stipulated time. The designated relay WT immediately acknowledges the receipt of the Search Lost WT Request message. It then refers to its network topology map information to steer its antenna beam in the direction of the lost WT, and sends a Search message to the lost WT. If the lost WT is able to receive the Search message, it responds with a Search Response message, and infers that it is lost. Note that the packet-by-packet adaptive beamsteering capability (see Section 2.3) enables the lost WT to quickly steer its beam towards the candidate relay WT to receive the Search message and respond back, in case there is no blockage in the direction of the candidate relay WT.

Upon receiving the Search Response message from the lost WT, the chosen relay WT sends a Search Outcome message to the AP, reporting successful lost node discovery and the quality of the link (i.e., the received signal strength) between itself and the lost WT. Otherwise, after waiting for a PIFS interval, the

chosen relay WT informs the AP of the lost node discovery failure via the Search Outcome message. Depending on the response from the designated relay WT, the AP decides whether to choose another WT to attempt the lost node discovery procedure or to use the current chosen WT as a relay for the future data transfers to the lost WT. Upon a successful lost WT discovery, the AP adds the required data transfer time for the lost WT to the relay WT's dwell time. The relay WT transparently interfaces the AP and the WT by forwarding the (data/ACK) MAC frames between them based on the destination MAC address of the frames, until the direct link is out of blockage.

The AP continues to send poll messages in the direction of the lost WT over the next superframes, in order to check if the direct link is restored. Once the obstruction is removed and the lost WT starts receiving direct transmissions from the AP, it responds to the AP's poll message. The AP switches back to the normal mode of operation after informing the relay WT to return to its previous state. The dwell times for the directional transmissions to the WTs are adjusted accordingly.

2.4.5 Achievable Rates

In this section, we estimate the aggregate data transfer capacity for our multi-hop MAC framework. We first find the aggregate throughput for the case when no

Parameter	Symbol	Value
PHY data rate	R	2 Gbps
Propagation delay	δ_p	50ns
PHY overhead	T_{PHY}	250ns
Header overhead (IP+UDP+MAC)	T_{hdr}	$56*8/R$
Payload Tx time	$T_{payload}$	$1000*8/R$
Short MAC frame Tx time	T_{ShFr}	$T_{PHY} + 14 * 8/R + \delta_p$
SIFS interval	T_{SIFS}	100ns
ACK Tx time	T_{ACK}	T_{ShFr}
TXOP duration	T_{TXOP}	100 μ s
Polling overhead	T_{poll}	$2T_{ShFr} + T_{SIFS}$
Maximum Trailing Control period	T_{TrCP}	50 μ s
Hello/Hello Response Tx time	T_H/T_{HR}	T_{ShFr}

Table 2.2: Protocol parameters

WT is blocked and packet transmissions to the WTs from the AP are single hop. Then we calculate the change in throughput for the cases when the LOS connectivity to some WTs is lost, and multihop relay is used as an alternative mechanism for data transfer. The underlying goal is to verify that multihop relay results in graceful degradation of overall network throughput rather than loss of connectivity to such WTs altogether, which is extremely undesirable. Given the high data rates afforded by 60 GHz transceiver systems, this reduction in the aggregate throughput does not affect the applications unless the network is operating at full capacity. Also, in order to provide the required quality of service (QoS) support to different WPAN applications, continuous connectivity to all the nodes is essential.

We consider QPSK modulation at 2 Gbps, as in the system described in Section 2.3. We assume that 200 symbols are required for beamformer training and

signal acquisition/synchronization, adding 200ns of overhead. While this is adequate for the single carrier system we envision, OFDM might incur additional overhead; the framework for our analysis, however, would remain identical. The physical layer control protocol overhead is assumed to be 50ns. Thus, the total PHY overhead for each data transmission is 250ns. Note that at the nanosecond scale, the propagation delays of signals can no longer be neglected. We assume a maximum propagation delay of 50ns. We also assume that all the nodes operate at the same data rate. The maximum allowed TXOP duration is assumed to be 100 μ s. The AP polls each registered WT once every superframe and checks connectivity. The superframe duration is allowed to vary as per the data requirements of WTs, but it is limited to a maximum duration (determined by the TXOP interval, the number of WTs, and the system configuration).

We find the maximum aggregate throughput sustained by the network, assuming backlogged UDP flows from the AP to all the WTs with a packet size (P_{size}) of 1000 bytes. This UDP based application model can incorporate the requirements of the common WPAN applications including streaming content download for HDTV, real time streaming and wireless data bus. Table 2.2 lists our notation and parameter values.

The total time required to transmit a packet is given by

$$T_{pkt} = T_{PHY} + T_{hdr} + T_{payload} + \delta_p = 4.524\mu s. \quad (2.5)$$

Every successfully received packet from the AP is acknowledged by an ACK message from the WT since the WTs do not have reverse data traffic. Thus, the total number of packets transmitted in one TXOP is given by

$$N_{pkt} = \frac{T_{TXOP} - T_{poll}}{T_{pkt} + T_{SIFS} + T_{ACK}} = 19. \quad (2.6)$$

In a network consisting of n active WTs fully utilizing their TXOP duration, the total superframe time will be $T_{SF} = n \cdot T_{TXOP} + T_{TrCP}$, which for an eight WT network evaluates to $850\mu s$. We note that this duration is still small enough such that the AP and WTs can closely monitor the changes in network topology and they can adapt to the dynamics of the indoor environment by using multihop relays. Thus, the aggregate throughput S_{SH} sustained by the network under consideration is given by

$$S_{SH} = \frac{n \cdot N_{pkt} \cdot P_{size}}{T_{SF}}, \quad (2.7)$$

which equals 1.43 Gbps for the eight WT example. If we assume that m out of n WTs in the network are connected through two hop paths, the T_{SF} increases by mT_{TXOP} to facilitate data transfer over multihop paths without any packet loss.

The aggregate throughput S_{MHR} in that case will be

$$S_{MHR} = \frac{n \cdot N_{pkt} \cdot P_{size}}{T_{SF} + m \cdot T_{TXOP}}. \quad (2.8)$$

Fig. 2.7 illustrates the aggregate network throughput for an eight WT network as a function of the number of blocked WTs using relays for data transfer. The re-

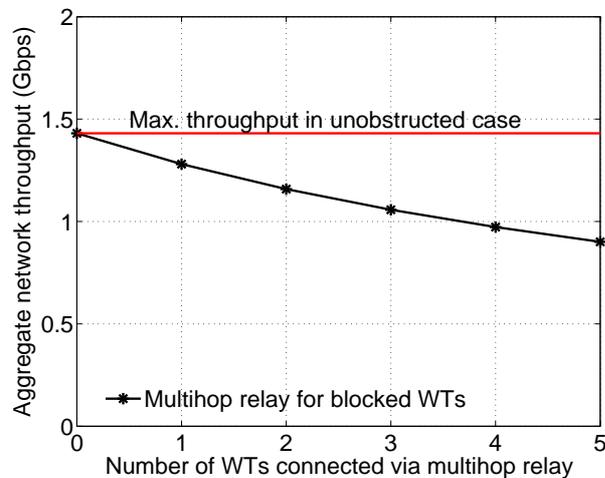


Figure 2.7: Aggregate network throughput for an eight WT network as a function of the number of blocked WTs connected via relays.

duction in aggregate network throughput (as compared with 1.43 Gbps in full single hop connectivity case) is the cost incurred in order to maintain connectivity to all the active WTs and to ensure that even the *lost* WTs do not suffer from packet loss because of blockage of the direct link to the AP. We therefore consider this reduction in the aggregate network throughput as graceful throughput degradation.

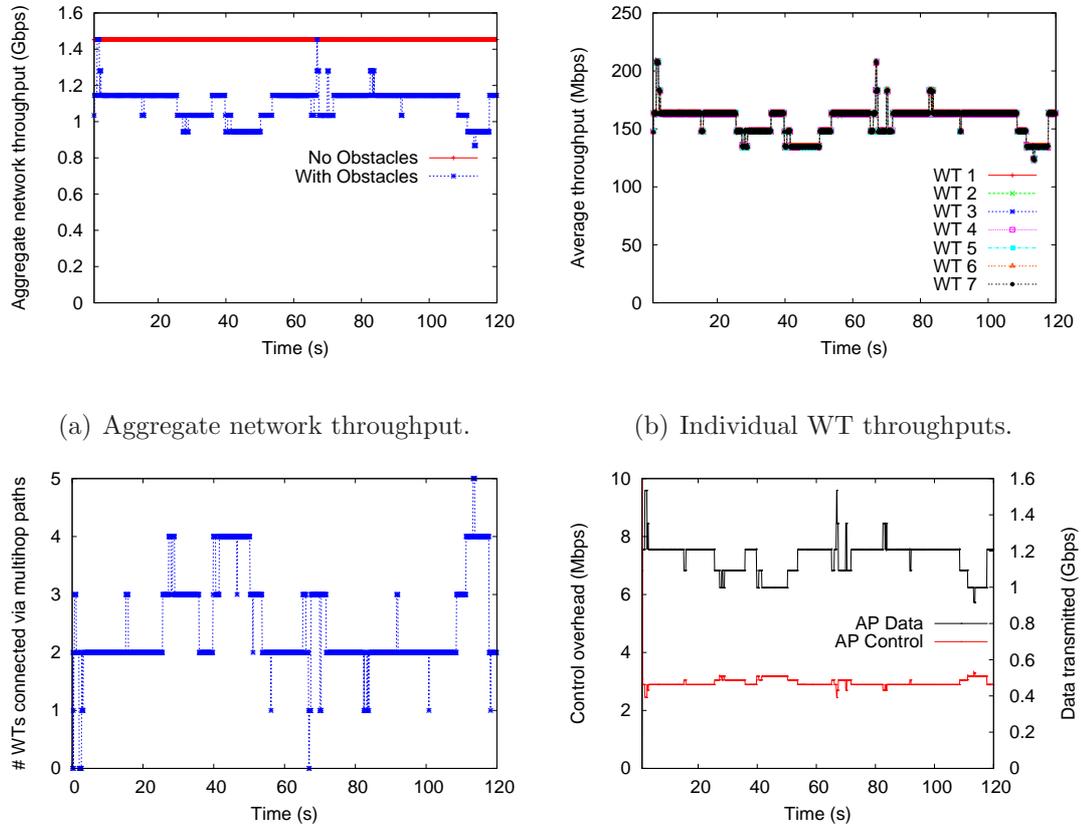
We note that at nanoseconds time scales, the packet processing delays at the AP and the WTs are non-negligible and should be accounted for in order to get the actual delays. Because these factors appear in all our time approximations, they will cause a similar shift in the estimates.

2.5 Performance Evaluation

2.5.1 Simulation Model

We first describe our simulation model which consists of the Matlab radio propagation tool that calculates the link losses (see Section 2.3), and the packet level simulations for performance evaluation of the directional relay MAC protocol. We then present the simulation results and discuss the insights obtained from the results.

Packet Level Network Simulation: We have implemented the multihop relay directional MAC protocol over the QualNet Network Simulator [61]. The link loss data obtained from the Matlab mm-wave propagation tool is fed into the physical layer module of the QualNet simulator, modified to model mm wave communication. The physical layer model also accounts for the antenna directivities while evaluating the signal to interference and noise ratio (SINR) values at each node, which are used to calculate the bit error rates (BER) based on the selected modulation. Packet loss probabilities are evaluated using the BER value and the packet size, which then determine successful or failed packet reception. The MAC and the PHY layer parameters used for the simulations are listed in Table 2.2. To obtain the maximum supported data rates in the system, each WT is assumed to



(a) Aggregate network throughput. (b) Individual WT throughputs. (c) Number of WTs connected via multihop paths. (d) AP control overhead and data transmission.

Figure 2.8: Protocol performance: living room scenario.

download UDP data from the AP such that the AP is always backlogged. The indoor environment is characterized by the parameters listed in Table 2.1. We refer back to the two example scenarios: the living room and the office space (see Figs. 2.3(a) and 2.3(b)) introduced in Section 2.3 to evaluate the performance of the multihop relay directional MAC protocol.

2.5.2 Simulation Results

Fig. 2.8(a) shows the aggregate network throughput as a function of time for the living room scenario - both for the case with obstacles and in an unobstructed environment. We observe that the aggregate throughput remains fairly consistent, even in the presence of the moving and stationary obstacles. The difference between the two curves quantifies the associated throughput drop (as compared with the throughput achieved in an unobstructed environment) when multihop relay is used to communicate with the blocked WTs that would have otherwise been unreachable from the AP.

The throughput achieved by the individual WTs in the living room setting is shown in Fig. 2.8(b). All the WTs effectively receive the same average throughput over time because of the equal service weights assigned to all the WTs in the simulations. The assignment of different transmission scheduling weights to packets based on the QoS requirements of the applications at each WT is a trivial extension of the equal weight case presented here. Fig. 2.8(c) presents the number of WTs connected via multihop paths at different sampling instances of the simulation. We infer that at any time instant, there are a significant number of WTs using multihop relay, which indicates the importance of multihop paths in maintaining uninterrupted network connectivity even under the scenarios where many LOS links are blocked because of obstacles.

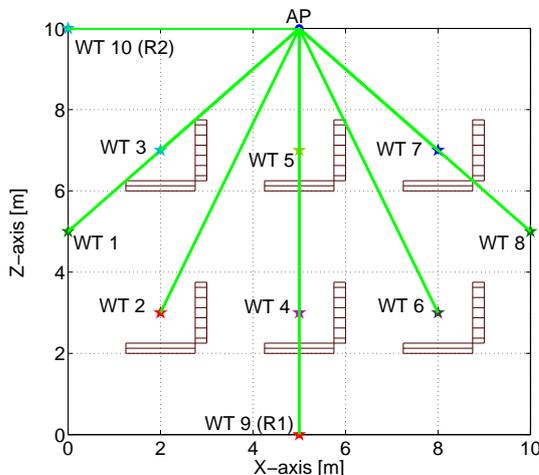
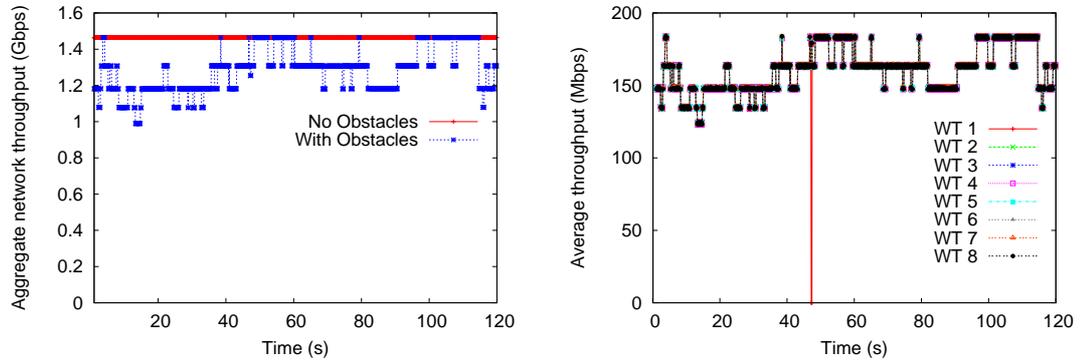


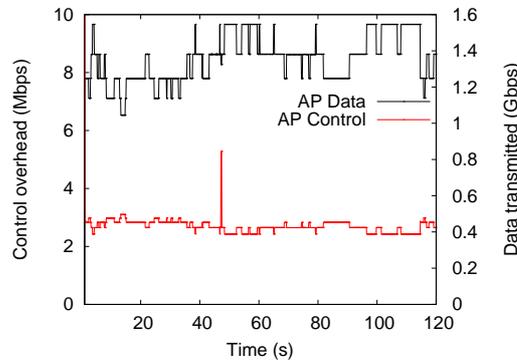
Figure 2.9: Office scenario with two relay WTs.

Fig. 2.8(d) illustrates the AP's data transmission rate and the control overhead due to polling, lost node discovery, and relay path setup, as a function of time for the living room scenario. The graph demonstrates the low relative control overhead, and also provides insight into the variation of the control message overhead as a function of the number of WTs on multihop connectivity. The direct relation of the control overhead and the number of nodes on multihop connectivity arises from the lost node discovery and the poll messages for the blocked WTs in each superframe. The AP needs to check whether the WT is back on LOS connectivity in every superframe. Although the AP remembers the corresponding relay WT used in the last superframe to avoid having to search for a relay in every superframe, it needs to check whether the previous relay node is still connected in the current su-



(a) Aggregate network throughput.

(b) Individual WT throughputs.

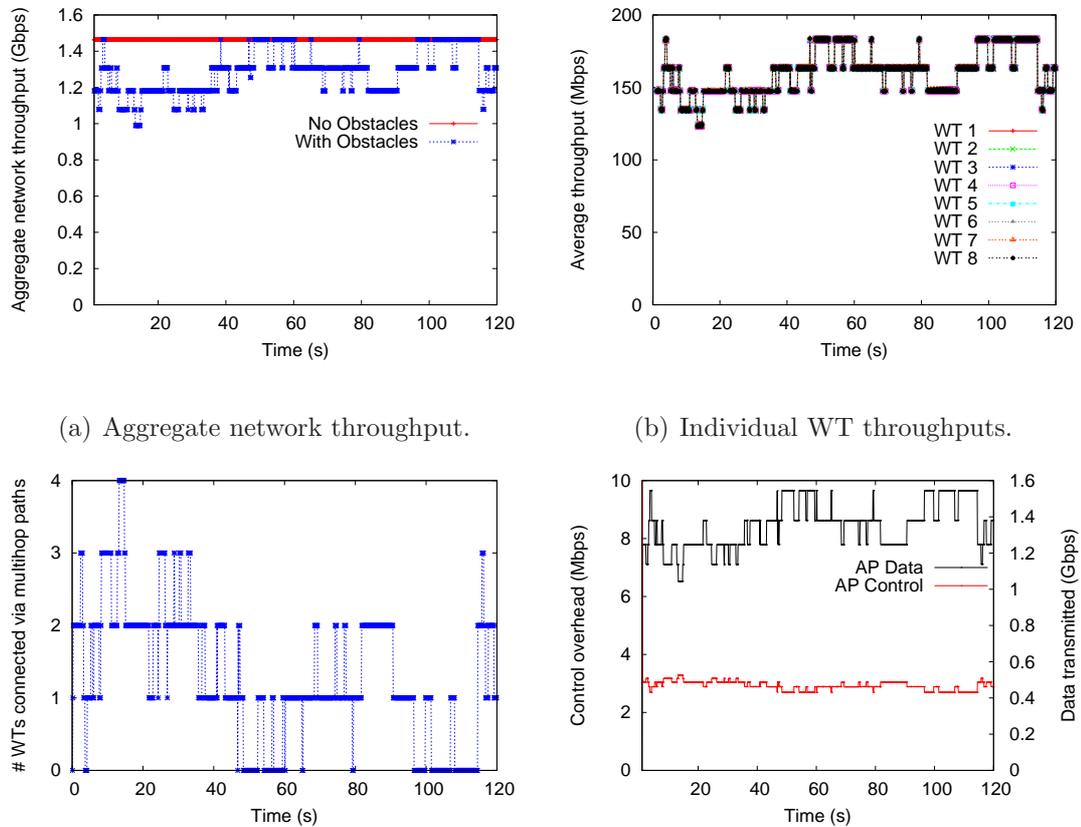


(c) AP control overhead and data transmission.

Figure 2.10: Protocol performance: office scenario.

perframe. This accounts for the proportional (although relatively small) increase in the control overhead with increasing number of WTs on multihop paths.

The aggregate network throughput for the office setting is plotted in Fig. 2.10(a). Fig. 2.10(b) shows the throughput achieved by the individual WTs. Fig. 2.10(b) presents an interesting scenario where WT1 is completely inaccessible (i.e., it is blocked from all the other nodes in the network) from 47.0 seconds to 47.3 seconds.



(a) Aggregate network throughput. (b) Individual WT throughputs. (c) Number of WTs connected via multihop paths (d) AP control overhead and data transmission.

Figure 2.11: Protocol performance: office scenario with two relay WTs.

Therefore, it achieves zero throughput because the lost node discovery procedure does not yield a relay node for multihop communication to WT1. This example demonstrates that the number and placement of relays play an important role in ensuring that the probability of complete blockage of the WTs is minimal. Now consider the scenario shown in Fig. 2.9 where an additional relay node is added

to the office network to ensure that no WT is completely isolated. Figs. 2.11(a) and 2.11(b) show the aggregate network throughput and the per-WT throughput for this setting, respectively. The additional relay adds redundancy by sharing the relay transmission load and ensures that the network is connected throughout the simulation duration. However, it is important to realize that 100% connectivity consistency cannot be guaranteed for all scenarios - one can imagine different pathological cases where a WT can be completely isolated, although a majority of these scenarios do not correspond to realistic use cases of WPAN devices. Usually the number of required relays depends on how challenging the indoor setting is in terms of the chance of complete blockage of WTs. This likelihood can be minimized by intelligently choosing relay positions (high up on the walls, or on the ceiling) such that they have a better chance of always being connected to the WTs.

Fig. 2.11(c) shows the number of WTs connected via multihop paths at different sampling instances of the simulation for the office scenario with two relay WTs. The inference is the same as for the living room scenario: multihop relay plays a critical role in maintaining network connectivity and high data rates to each WT.

Figs. 2.10(c) and 2.11(d) illustrate the AP's data transmission rate and the associated control overhead as a function of time for the two office scenarios.

Fig. 2.10(c) provides an interesting insight. The spike in the control overhead for the office scenario with a single relay at around 47 seconds is because of the AP's repeated lost node discovery attempts to search for the lost WT1 that is completely isolated for 0.3 seconds. In this example, the AP is configured to continue searching for the lost WT to illustrate the possible high overhead in such cases because of the repeated lost node discovery attempts. This spike can easily be avoided by configuring the AP to search for a lost WT only for a short period of time (equivalently, the AP performs lost node discovery attempts only over a limited number of contiguous superframes), after which the AP should consider the WT to have left the network. The lost WT can join the network again when it comes out of complete blockage (i.e., it can reach one of the network WTs or the AP). Figs. 2.10(c) and 2.11(d) also reaffirm the insight on the relationship between the number of nodes on multihop connectivity and the control overhead obtained from the living room scenario.

QoS Performance: Since the multihop directional relay MAC protocol is essentially a contention free transmission scheduling-based MAC protocol, it can offer assured QoS to different applications over the typical WPAN scenarios. Moreover, the protocol can easily be extended to incorporate the well known QoS enhancing MAC features such as block/no acknowledgments and traffic prioritization in IEEE 802.11e HCCA [62]. Since our motivation is to illustrate the promise of mul-

tihop relay and directional communication as key enablers for 60 GHz WPANs, we focus on showing their potential via our base protocol.

2.6 Summary

Our results illustrate the critical role of cross-layer design in exploiting the large unlicensed bandwidth available in the 60 GHz band. The diffraction-based connectivity model is an effective tool for cross-layer design: it yields results that conform to our intuition that directional LOS mm wave links experience relatively high levels of outage due to stationary and moving obstacles. Despite this fragility of the individual mm wave links, we show via extensive packet-level simulations that the proposed multihop MAC architecture is successful in providing robust connectivity in typical “Superbowl Party” and office settings. Thus, unlike the infrastructure mode operation in the 2.4 GHz and 5 GHz WLANs where the WTs communicate with the AP over a single hop, we believe that multihop communication, possibly with nodes explicitly designated as relays, must play a fundamental role in 60 GHz WPANs.

Chapter 3

60 GHz Outdoor Mesh Networks

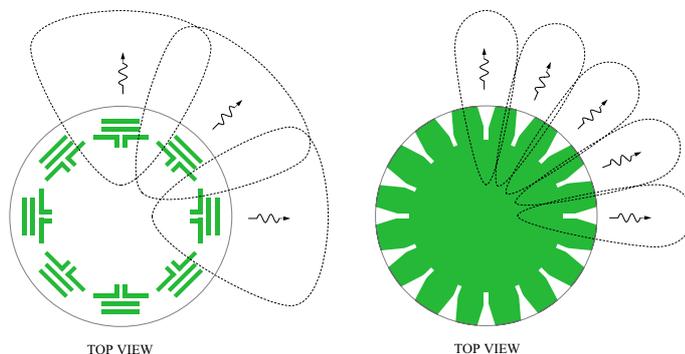
While recent progress in RFIC design in the mm-wave band has spurred significant recent interest in 60 GHz networks for indoor multimedia applications, our focus in this chapter is on multiGigabit *outdoor* 60 GHz mesh networks with link ranges on the order of 100 meters. Such mm wave mesh networks can provide a high-speed backhaul needed for broadband connectivity in regions where there is limited wired or optical infrastructure. Further, these networks can also serve as a means of stimulating competition by enabling smaller broadband service providers to bypass the existing infrastructure provided by entrenched carriers.

3.1 Introduction

In this chapter, we take the first steps towards defining an architecture for outdoor 60 GHz mesh networks, accounting for the unique characteristics of mm wave communication. Specifically, our goal is to address what we believe is a

critical bottleneck in the design of 60 GHz mesh networks: medium access control (MAC) for highly directional links. We constrain the links to be highly directional because it is both *essential* and *feasible*. Directionality is *essential* because, for omnidirectional transmission, the path loss scales as λ^2 , where λ is the carrier wavelength, so that 60 GHz is 22 dB worse than 5 GHz. Given the difficulty of producing RF power at mm wave frequencies, such a loss in the link budget is unacceptable. By using narrow beams at both the transmitter and receiver, however, we can more than compensate for the loss: for the same antenna size, the directivity scales as λ^{-2} , so that the overall path gain with directive antennas at both ends scales as λ^{-2} , leading to a net *gain* of 22 dB over 5 GHz. Directionality is *feasible* because, at these small wavelengths, large-scale antenna arrays can be realized as patterns of metal on circuit board, and can be employed to synthesize electronically steerable (e.g., see [63,64]), highly directional beams for nodes with form factors comparable to that of a WiFi access point.

The use of highly directional links requires a complete rethinking of MAC protocol design relative to what we are used to in CSMA-based WiFi networks for two reasons. Since spatial interference is significantly reduced, interference management or avoidance becomes a secondary consideration. On the other hand, the deafness resulting from directionality means that carrier sensing or promiscuous listening in order to monitor the state of one's neighbors is no longer feasible.



(a) Yagi-Uda antenna array (b) Open-slot antenna array

Figure 3.1: Steerable arrays with high directivity and compact form factors.

In this chapter, we first use a statistical analysis of interference to show that, for the antenna directivities of interest to us, even uncoordinated transmission for different transmit-receive pairs leads to small collision probabilities. This motivates a *pseudo-wired* abstraction to serve as a first order approximation of mm-wave wireless links, and provides insights to guide MAC design. Under this abstraction, transmissions on different links do not interfere with each other. However, unlike a truly wired node, mm-wave network nodes have a half-duplex constraint, i.e., they can only send or receive at a given time.

Guided by these insights on mm wave links, the second challenge that we take up is whether we can design a *distributed* MAC protocol that is scalable and robust and that achieves coordination in the face of deafness in order to efficiently utilize the medium.

Before discussing network design, we note that electronically steerable links can be realized for easy-to-deploy mesh nodes that can connect with a neighbor node regardless of the direction of the neighbor's location. Several concept designs for steerable arrays with a 360° field of view for such *omni-coverage yet highly directional* nodes are currently being explored: Figs. 3.1(a) and 3.1(b) depict two representative array implementations with Yagi-Uda and open-slot antennas, respectively. Each element in the array can have significant directivity (e.g., up to 20 dBi), so that such arrays can be operated in beamswitched mode, but even higher directivity, up to 26 dBi, can be obtained by combining signals from a small number of elements with overlapping fields of view. While hardware optimizations for efficient realization of baseband and RF processing at such nodes are topics of ongoing research, the hardware developed for emerging indoor 60 GHz products [4,5] imply that there are no fundamental roadblocks in realizing such nodes.

Contributions: We first provide a probabilistic analysis of the interference incurred due to uncoordinated transmissions in a 60 GHz outdoor mesh network, and show that, for the parameters considered, the links in the network can be thought of as *pseudo-wired*. That is, interference can essentially be ignored in MAC design, and the challenge is to schedule half-duplex transmissions in the face of the “deafness” resulting from highly directional links.

We next present a *distributed* MAC protocol that attains high medium utilization, comparable to efficient time division multiplexing (TDM), while avoiding the network-wide explicit coordination required to implement global TDM schedules. The key idea is to use memory to achieve implicit coordination among the mesh nodes despite deafness. Each node persists in using the transmit or receive slots that have been successfully used for a given neighbor, and puts slots over which it has been unsuccessful on a per-neighbor blacklist. There is no attempt at proactive interference avoidance and the blacklisting is oblivious to whether the failure was due to deafness (more likely) or interference (less likely). Thus, a node’s transmit and receive history with each of its active neighbors provides feedback that is used for implicit coordination, and persistent use of a given slot for transmitting to a given neighbor leads to an approximate TDM schedule. The system is prevented from locking into undesirable schedules by each node leaving “enough room” in its schedule, along with suitable randomization of persistence and blacklist lifetimes, for adaptively accommodating changes in demand.

The novelty of our approach to MAC design lies in the use of simple learning rules employed by each mesh node, based *only* on the memory of its own transmit/receive outcomes, to converge to TDM-like schedules in a distributed fashion with minimal control overhead. In addition to satisfying the design constraints unique to highly directional 60 GHz networks, this approach obviates the need

for frequent control message exchanges for coordination or local neighborhood information gathering/dissemination typical of the prior works on distributed MAC scheduling. The advantages over a centralized approach are also worth stating: lower overhead, and better scalability and fault tolerance.

We offer approximate analytical insights into the working of our protocol via a Markov chain fixed-point analysis, and study the properties of the approximate TDM schedules that arise from our protocol. We present an extensive evaluation of our protocol using the QualNet simulator [61], modified to model 60 GHz communication.

3.2 Related Work

To the best of our knowledge, this is the first work on interference analysis and distributed MAC protocol design for mm wave mesh networks.

Directional networking for multihop wireless networks in the WiFi band has been extensively studied in the past [23,24,27,28,65–67]. Many of these proposals employ directional networking as an option to enhance the network performance (via high spatial reuse or range extension) while still exploiting the broadcast nature of omnidirectional communication for critical control message propagation. Thus, such protocols do not apply to 60 GHz mesh networks where there is no

omnidirectional mode. Further, most directional protocols designed for lower frequency bands consider interference management as a prominent MAC design goal and use explicit signaling for proactive interference avoidance, since the achievable antenna directivities are much smaller than those possible at 60 GHz.

A few protocols that employ an out-of-band tone to compensate for the loss of coordination [27, 66, 67] require additional control and hardware complexity: [66, 67] need multiple simultaneously active transceivers to concurrently transmit a directional out-of-band busy-tone, whereas [27] requires delayed transmission of an omnidirectional out-of-band tone. A variation of slotted Aloha is proposed in [28], where each idle node equipped with an adaptive antenna array beamforms towards the strongest tone (tones precede every packet transmission), and forms nulls towards potential interferers. Another approach is to use directional flooding of control messages on all sectors [24]; this would have prohibitive overhead for 60 GHz nodes employing a large number of narrow sectors.

While our goal here is to design a MAC protocol providing high network utilization and a rough measure of fairness, the pseudowired abstraction presented here can also form the basis for scheduling algorithms (perhaps layered as a small perturbation on top of a MAC such as ours) that explicitly account for traffic patterns. This is an important topic for future work, especially for providing the QoS that may be required for a wireless backhaul network. We anticipate that

the large body of work on distributed link scheduling can be leveraged for this purpose [68–70].

Note that we have previously employed the idea of “sticking” to TDM-like schedules in a completely different context: efficient support of periodic traffic in omnidirectional CSMA networks [9]. In [9], transmissions stick to periodic schedules after a successful initial handshake, and use the memory of carrier-sensed activity to avoid contention. Related work includes [71], which uses two priority level broadcast control messaging over each slot, guided by the memory of the past transmission outcomes. While the application of memory in this chapter is quite different (enabling coordination in the face of deafness), all these works indicate that memory and learning could potentially be powerful (and presently underused) tools in network protocol design.

3.3 60 GHz Outdoor Mesh Architecture

We envision a rooftop or lamp-post based network with line of sight (LoS) links. Our baseline scenario, used throughout this chapter, is a 2 Gbps link at a range of 100 meters. Assuming Quadrature Phase Shift Keying (QPSK) signaling (2 bits/symbol) at 1 Gsymbol/sec over a bandwidth of 1.5 GHz, for a desired signal to noise ratio (SNR) of 15 dB (which allows for uncoded error probability

of 10^{-9} or better), Oxygen absorption loss at 15 dB/km, and 10 mW transmit power, we need antenna gains of about 24 dBi at both the transmitter and the receiver in order to provide a 10 dB link margin.

In addition to the half duplex constraint, we assume that each node can only communicate with a single neighbor during a time slot: all the other signals are treated as interference. Techniques such as multiuser detection and interference nulling could potentially enhance network performance, but our objective here is to explore the performance with a simple physical layer implementable in the short term.

While we focus on MAC design here because it presents the greatest challenge, requiring adaptation at the packet time scale, we comment briefly on possible approaches for slower time scale mechanisms such as network discovery and slot-level synchronization, which the MAC protocol builds on. We assume that nodes obtain information about their neighbors as well as the antenna array weights to communicate with them via a neighbor discovery procedure executed during network initialization. For example, a distributed variant of the algorithm proposed in [6] can be employed: a discovering node sequentially scans each sector (angular region determined by the antenna beamwidths) for the presence of neighbors by sending a hello message and waits for the response(s). The hello message payload is preceded by a preamble that acts as a training sequence for the receiving nodes to

beamform in the direction of the transmit node. The beamformed antenna array weights obtained by the receiving nodes are recorded and used while responding back, using reciprocity. Unregistered neighbors that successfully receive a hello message respond within a time window, with a random transmit start time delay to avoid a collision at the discovering node. This procedure is repeated multiple times such that all neighbors are registered, where the hello message sent over the last round requires any unregistered node to respond immediately to verify that no neighbor is left unregistered. The order in which nodes perform network discovery can either be coordinated by the gateway nodes, or it can be configured by the network administrator. Note that the execution of the discovery procedure for the complete network does not take more than a few seconds, which is acceptable for bootstrapping the network.

We assume a time-slotted system where each node is coarsely synchronized with its immediate neighbors, i.e., the slot boundaries are assumed to be aligned to within a guard time. We define a frame to be a sequence of k contiguous time slots. While the frame duration is assumed to be common, the frame start times or the time-slot indexing may be different for different nodes. In principle, a coarse time synchronization can initially be achieved during network discovery and re-calibrated thereafter, via the gateway nodes. Note that for outdoor mesh nodes, GPS-based synchronization is also a viable option. While a detailed in-

Investigation of distributed or gateway-led synchronization mechanisms is beyond the scope of this chapter, we note that effective local synchronization mechanisms that account for the high link directionality can be devised, for example, using distributed approaches inspired from [72, 73].

We assume that after network initialization, a receiving node has the ability to “tune into” a transmission from any of its neighbors. Such capability, which greatly simplifies MAC design, does not necessarily require advanced signal processing. For example, a receiver can rapidly scan its antenna beam in the (known) directions of all of its neighbors at the beginning of a time slot to look for possible incoming transmissions. In this manner, the signal acquisition and beamforming overhead can be kept reasonably low.

A complete network architecture also requires design and optimization of the higher layers, including possibly admission control, routing, and traffic-sensitive scheduling, in addition to MAC. A detailed investigation of these issues is beyond the scope of this chapter. In our MAC performance evaluations, we restrict attention to two scenarios: saturated traffic (to evaluate medium utilization efficiency) and constant bit rate traffic with static routing (to evaluate delay and delay jitter performance).

3.4 Network Model

For our interference analysis, we consider a Poisson distribution of nodes over a large area with a density ρ_s . If we now randomly select a subset of N_T nodes as transmitters, the distribution of transmitters on the area of interest is also Poisson, with density $\rho = \rho_s p_t$ where p_t is the probability that the selected node is transmitting. For a large deployment area A , we have $\rho \approx \frac{N_T}{A}$.

We assume that a transmission is successfully decoded by the receiver if the total signal to interference and noise ratio (SINR) is above a given threshold, say $\beta = 15$ dB (which allows for uncoded QPSK modulation at a BER of $< 10^{-9}$). Otherwise, a *collision* occurs and the transmission is lost. The amount of interference depends on the location of the interferer(s) relative to the receiver, and the radiation patterns of the antennas at the receiver and the interferer(s).

The standard Friis transmission equation gives the received power as a function of range r as

$$P_R(r) = P_T G_R G_T \left(\frac{\lambda}{4\pi r} \right)^2 e^{-\alpha r} \quad (3.1)$$

where P_T is the transmitted power, G_R , G_T are the gains of the receive and transmit antennas respectively, λ is the wavelength, and α is the attenuation factor due to absorption in the medium. For a mm wave link at 60 GHz, $\lambda = 5$ mm and α can be as high as 15 dB/km. Since lower absorption rates leads to

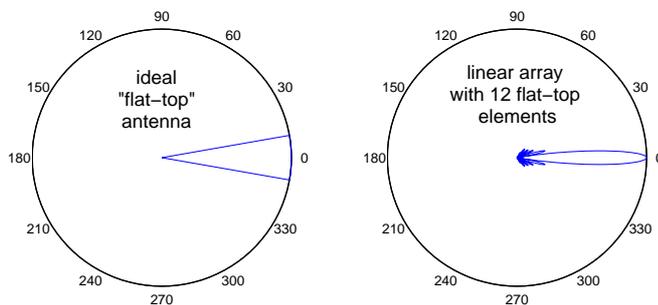


Figure 3.2: Gain pattern for a flat-top antenna and a linear array of flat-top elements.

more interference, we use the conservative value of $\alpha = 10$ dB/km in our numerical results for interference analysis.

3.4.1 Directional Antenna Model

Directional antennas are characterized by their pattern functions that measure the power gain $G(\phi, \theta)$ over the spherical azimuthal and elevation angle coordinates ϕ , θ . We assume that all nodes are on the same horizontal plane, and do not consider variation in beam pattern over the elevation angle θ , and work with the normalized 2-dimensional pattern

$$g(\phi) \triangleq \frac{G(\phi, 0)}{G_{max}} \text{ where } G_{max} = \max_{\phi} G(\phi, 0) \quad (3.2)$$

The azimuthal beamwidth of the antenna is then given by

$$\Delta\phi = \int_{\phi=-\pi}^{\pi} g(\phi) d\phi \quad (3.3)$$

One idealization that proves to be very useful for our interference analysis is that of a sectorized “flat-top” directional antenna, which has unit gain within its beamwidth and zero gain outside. More precisely,

$$g(\phi) = \begin{cases} 1, & |\phi| \leq \frac{\Delta\phi}{2}, \\ 0, & \text{otherwise} \end{cases} \quad (3.4)$$

While the flat-top antenna is an useful idealization, practical directional antenna gains have a more complex dependence on the azimuth angle. For instance, sidelobes in the gain function could cause significant interference even in directions far from the antenna boresight. While exact computation of the gain functions of practical mm-wave antenna arrays can be messy (because array elements usually are directional themselves), we can obtain useful models with some simplifying assumptions. In Fig. 3.1 only a subset of the array elements with significantly overlapping mainlobes contribute to the antenna gain in any specific direction. The number of elements in this subset can still be quite large: for instance if the diameter of the overall array is 10 cm, and the inter-element spacing is a half-wavelength, and each element has a azimuthal beamwidth of, say, 120° , we can have on the order of 20 – 30 elements with overlapping beams. If we assume that the beams of these elements are approximately flat, and neglect the curvature of their placement, we can approximate this sub-array as a uniform linear array, each of whose elements has a flat-top response. In particular, we obtain

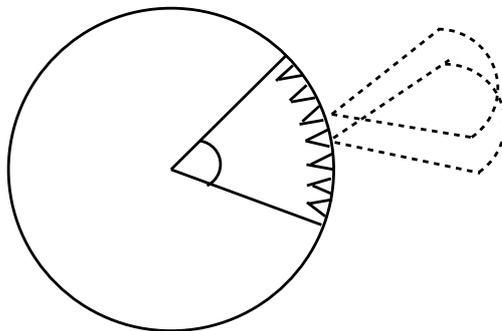


Figure 3.3: Approximating a circular array of slot antennas as a uniform linear array of flat-top elements.

the following gain function for an N element linear array in which each individual flat-top element has beamwidth $\Delta\phi_0$:

$$g(\phi) = \begin{cases} \frac{1}{N} \frac{\sin(\frac{N}{2}\pi \sin \phi)}{\sin(\frac{1}{2}\pi \sin \phi)}, & |\phi| \leq \frac{\Delta\phi_0}{2}, \\ 0, & \text{otherwise} \end{cases} \quad (3.5)$$

Fig. 3.2 shows the beam patterns for a narrow beam flat-top antenna and a 12-element linear array of broad-beam flat-top elements. The beamwidth in both cases (as defined as in (3.3)) is the same, 20° . While we have chosen a 20° beamwidth for ease of display, in our numerical results, we use a beamwidth of 10° ; this requires a 27-element array of flat-top antennas with sector-width $\Delta\phi_0 = 120^\circ$. Fig. 3.3 illustrates how to approximate a circular array with a linear array with flat-top elements.

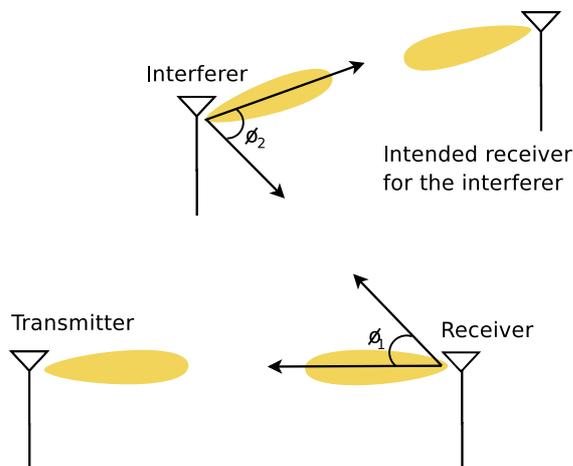


Figure 3.4: The geometry of interference with directional antennas.

3.5 Interference Analysis

We now investigate the validity of a pseudo-wired model for the links in a mm wave mesh network, by analyzing the probability of packet failure for uncoordinated transmissions. Consider the transmitter-receiver pair shown in Fig. 3.4; without loss of generality, assume that the receiver is located at the origin and is communicating with the transmitter located along the X-axis, at a distance less than or equal to the reference link distance R_0 while undergoing interference from other concurrent transmissions. The other $N_T - 1$ interfering transmitters are randomly placed over the area A , and are transmitting to receivers located at randomly chosen orientations.

In their well-known work on wireless network capacity [1], Gupta et al introduce two different models of interference. In the *protocol model*, a packet loss

occurs if and only if there is some interfering node whose signal at the receiver exceeds a given threshold. In the *physical model*, a packet loss occurs when the *total* interference from all nodes exceeds a given threshold. The models are summarized as follows:

$$\Pr(\text{collision}) \triangleq \begin{cases} \Pr(\max_k P_k \geq \frac{1}{\beta} P_0), & \text{(protocol model)} \\ \Pr(\sum_{k=1}^{N_T-1} P_k \geq \frac{1}{\beta} P_0), & \text{(physical model)} \end{cases}$$

where P_k is the power at the receiver of the signal from the k 'th interferer. Since $\sum_k P_k \geq \max_k P_k$, it follows that the $\Pr(\text{collision})$ for the physical model is lower-bounded by the corresponding value for the protocol model. We first derive an expression for the collision probability under the protocol model for the flat-top antenna, and extend the derivation to more general directional antennas, and then to the physical model. Our derivation is similar to the analysis of localization error in [74].

3.5.1 Protocol Model with Ideal Flat-Top Antennas

For the ideal flat-top antenna, only interferers located within the boresight of the receiver can cause a collision. Further, a transmitting node within this sector causes interference only if the receiver is within its boresight, which has probability $q = \frac{\Delta\phi}{2\pi}$ (since the potentially interfering transmitter is sending to a randomly chosen receiver). Let R_i be the *interference range* i.e. the maximum

distance an interferer can be from the receiver and still cause a collision. Using (3.1), the signal and interference powers are evaluated as:

$$P_R \equiv P_0 = P_T G_{max}^2 \left(\frac{\lambda}{4\pi R_0} \right)^2 e^{-\alpha R_0} \quad (3.6)$$

$$P_{interf} = P_T G_{max}^2 \left(\frac{\lambda}{4\pi R_i} \right)^2 e^{-\alpha R_i} \quad (3.7)$$

where we used $G_T = G_R = G_{max}$ for the antenna gains assuming that the interferer and the receiver are within each other's boresights. We set $P_R = P_0$: when the transmitter and receiver are steered towards each other, this is the signal power designed for at the reference distance R_0 . Using the collision condition $P_{interf} \geq \frac{1}{\beta} P_0$, we can rewrite (3.7) as:

$$\frac{R_i^2}{R_0^2} e^{\alpha(R_i - R_0)} = \beta \quad (3.8)$$

which determines R_i as a function of the SINR threshold β .

The number of potentially interfering transmitters is therefore a Poisson random variable with mean ρA_i , where $A_i = \frac{1}{2} \Delta \phi R_i^2$. The probability of any of these actually causing a collision is q , so that the number of interferers N_i causing a collision is also Poisson, with mean $q\mu_i$. The probability of a collision is therefore

given by

$$\Pr(\text{collision}) \equiv \Pr(N_i > 0) = 1 - e^{-q\mu_i} \quad (3.9)$$

$$= 1 - e^{-\frac{(\Delta\phi)^2}{4\pi}\rho R_i^2} \equiv 1 - e^{-\beta\rho R_0^2 A_c},$$

$$\text{where } A_c \triangleq \frac{(\Delta\phi)^2}{4\pi} e^{-\alpha(R_i - R_0)} \quad (3.10)$$

For a beamwidth of $\Delta\phi = 10^\circ$, $\beta = 15$ dB and $\rho R_0^2 = 1$ (corresponding to roughly $\pi\rho R_0^2 \approx 3$ transmitting nodes within communication range of each receiver), (3.10) gives an estimate of $\Pr(\text{collision}) \approx 3.7\%$, which suggests that acceptable MAC performance may be possible with minimal coordination for interference management.

3.5.2 Protocol Model with General Directional Antennas

We now generalize (3.10) to a general directional antenna. We first compute the probability of collision due to a single interferer at a fixed location at a distance r , and angle ϕ_1 relative to the receiver as shown in Fig. 3.4. The angle ϕ_2 represents the direction of the interferer's beam relative to the receiver. We model ϕ_1, ϕ_2 as independent and uniformly distributed over $(-\pi, \pi]$, given the random orientation of the interfering transmitter and its beam relative to the desired receiver. Signal power is still given by (3.6), and the interference power is:

$$P_{interf} = P_T G_{max}^2 g(\phi_1) g(\phi_2) \left(\frac{\lambda}{4\pi r}\right)^2 e^{-\alpha r} \quad (3.11)$$

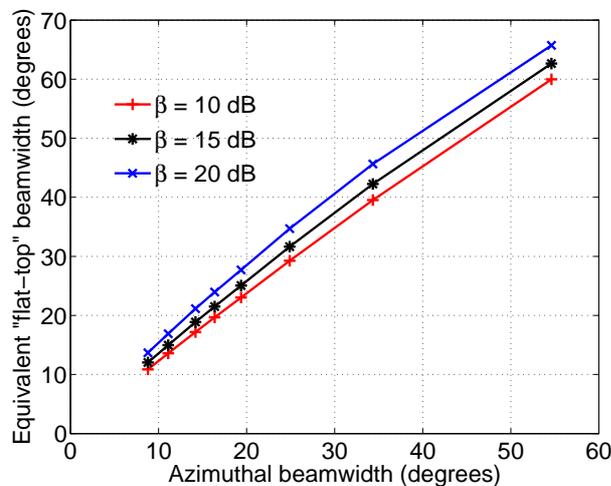


Figure 3.5: Flat-top beamwidth.

where we used $G_R = G_{max}g(\phi_1)$ and $G_T = G_{max}g(\phi_2)$. Using (3.6), we can rewrite

(3.11) as

$$P_{interf} = P_0 g(\phi_1) g(\phi_2) \left(\frac{R_0}{r} \right)^2 e^{-\alpha(r-R_0)} \quad (3.12)$$

Therefore the probability $p_c(r, \phi_1)$ that this interferer would cause a collision is

$$\begin{aligned} p_c(r, \phi_1) &\triangleq \Pr \left(P_{interf} \geq \frac{1}{\beta} P_0 \right) \\ &= \Pr \left(g(\phi_1) g(\phi_2) \geq \frac{1}{\beta} \left(\frac{r}{R_0} \right)^2 e^{\alpha(r-R_0)} \right) \\ &= \frac{1}{2\pi} \int_{-\pi}^{\pi} \mathbf{1} \left(g(\phi_1) g(\phi_2) \geq \frac{1}{\beta} \left(\frac{r}{R_0} \right)^2 e^{\alpha(r-R_0)} \right) d\phi_2 \end{aligned} \quad (3.13)$$

where $\mathbf{1}(\cdot)$ is the indicator function that takes the value 1 when its argument is true, and 0 otherwise.

Consider now an interferer placed at random within a large area A . Then the probability of collision \hat{p}_c can be obtained by averaging (3.13) over all possible

positions of the interferer:

$$\begin{aligned}\hat{p}_c &= \frac{1}{A} \int \int_{(r, \phi_1) \in A} p_c(r, \phi_1) r dr d\phi_1 \\ &= \frac{\beta R_0^2}{A} \int \int_{\hat{r}, \phi_1} p_c(\sqrt{\beta} R_0 \hat{r}, \phi_1) \hat{r} d\hat{r} d\phi_1\end{aligned}\quad (3.14)$$

where we set $\hat{r} \triangleq \frac{1}{\sqrt{\beta}} \frac{r}{R_0}$, i.e. \hat{r} is the distance of the interferer normalized to the reference link distance R_0 and interference threshold β . Using (3.13) in (3.14), we get

$$\begin{aligned}\hat{p}_c &= \frac{\beta R_0^2}{2\pi A} \int \int_{\hat{r}, \phi_1} \int_{\phi_2 = -\pi}^{\pi} \mathbf{1}\left(g(\phi_1)g(\phi_2) \geq \hat{r}^2 e^{\alpha R_0(\sqrt{\beta}\hat{r}-1)}\right) \\ &\quad d\phi_2 \hat{r} d\hat{r} d\phi_1\end{aligned}\quad (3.15)$$

We now let the area A become infinitely large and cover the whole plane. From (3.2), $g(\phi_1)$ and $g(\phi_2)$ are upper-bounded by 1. We therefore only need to consider $\hat{r} \leq e^{\frac{\alpha R_0}{2}}$ in (3.15), because the argument of the indicator function $\mathbf{1}(\cdot)$ in (3.15) is always false outside this range. Then we have

$$\hat{p}_c = \beta R_0^2 \frac{A_c}{A}, \text{ where} \quad (3.16)$$

$$\begin{aligned}A_c &\triangleq \frac{1}{2\pi} \int_{\hat{r}=0}^{e^{\frac{\alpha R_0}{2}}} \int \int_{\phi_1, \phi_2 = -\pi}^{\pi} \mathbf{1}\left(g(\phi_1)g(\phi_2) \geq \hat{r}^2 e^{\alpha R_0(\sqrt{\beta}\hat{r}-1)}\right) \\ &\quad d\phi_2 \hat{r} d\hat{r} d\phi_1\end{aligned}\quad (3.17)$$

We now consider $N_T - 1 = \rho A$ interferers placed randomly in the area A . Each interferer has a collision probability \hat{p}_c with the receiver given by (3.16). A collision

occurs if at least one of these interferers cause a collision, and its probability is given by

$$\begin{aligned} \Pr(\text{collision}) &= 1 - (1 - \hat{p}_c)^{N_T-1} = 1 - \lim_{A \rightarrow \infty} \left(1 - \beta R_0^2 \frac{A_c}{A}\right)^{\rho A} \\ &= 1 - e^{-\beta \rho R_0^2 A_c} \end{aligned} \quad (3.18)$$

Since (3.18) has an identical form to (3.10), the collision probability depends on the antenna pattern only through A_c . Thus, for the protocol model, we can restrict attention to an equivalent flat-top model whose beamwidth can be calculated from (3.10) and (3.16) as $\Delta\phi_{eq} \triangleq \sqrt{4\pi A_c} e^{\frac{\alpha}{2}(R_i - R_0)}$. For instance, a linear 24-element linear array of flat-top antennas of sector size 120° and half-wavelength spacing has an equivalent “flat-top” beamwidth of about 15° for $\alpha = 10$ dB/km. Fig. 3.5 shows the equivalent “flat-top” beamwidths for linear arrays of different numbers of flat-top elements and half-wavelength spacing; as seen from the figure, the “flat-top” beamwidth is numerically close to the algebraic beamwidth given by (3.3) and does not vary much with the SINR threshold β .

3.5.3 Physical Model

In the protocol model, only nodes located within a bounded distance from the receiver are capable of causing a collision. On the other hand, for the physical model, interfering signals from a large number of far-away transmitters could, in

principle, sum up at a desired receiver to cause packet failure. It turns out that traditional upper bound techniques such as Chernoff and Markov-type bounds do not work well when characterizing sum interference over a large area. We therefore use a hybrid approach, using an analytical Markov upper bound to characterize the effect of “far-away” interferers, and characterizing the sum interference from interferers within a bounded region through Monte-Carlo simulations. Let r_k be the distance of the k 'th interferer from the receiver. We write the total interference power as the sum of two contributions P_{near} and P_{far} , defined as:

$$P_{near} \triangleq \sum_{\{k:r_k \leq R_{th}\}} P_k, \text{ and } P_{far} \triangleq \sum_{\{k:r_k > R_{th}\}} P_k \quad (3.19)$$

where R_{th} is a suitable large distance, say $R_{th} = 40R_0$. Then we have

$$\begin{aligned} \Pr(\text{collision}) &= \Pr(P_{near} + P_{far} \geq \frac{1}{\beta}P_0) \\ &\leq \Pr(P_{far} \geq \frac{1}{\beta}\Delta P) + \Pr(P_{near} \geq \frac{1}{\beta}(P_0 - \Delta P)) \\ &\leq \frac{\mathbb{E}[P_{far}]}{\Delta P/\beta} + \Pr(P_{near} \geq \frac{1}{\beta}(P_0 - \Delta P)) \end{aligned} \quad (3.20)$$

where we used the Markov Inequality to bound the first term in (3.20). The expectation in the first term is readily evaluated as:

$$\begin{aligned}
 \mathbb{E}[P_{far}] &\equiv \frac{\rho}{2\pi} \int_{\hat{r}=R_{th}}^{\infty} \int_{\phi_1, \phi_2=-\pi}^{\pi} P_0 g(\phi_1) g(\phi_2) \frac{R_0^2}{r^2} e^{-\alpha(r-R_0)} \\
 &\quad r dr d\phi_1 d\phi_2 \\
 &= \frac{(\Delta\phi)^2}{2\pi} P_0 (\rho R_0^2) e^{\alpha R_0} \int_{r=R_{th}}^{\infty} \frac{e^{-\alpha r}}{r} dr \\
 &\lesssim \frac{(\Delta\phi)^2}{2\pi} \frac{P_0 (\rho R_0^2)}{\alpha R_{th}} e^{\alpha(R_0 - R_{th})} \tag{3.21}
 \end{aligned}$$

3.5.4 A Pseudo-Wired Abstraction

Figs. 3.6 and 3.7 show the collision probabilities for ideal “flat-top” arrays and for linear arrays with $\rho R_0^2 = 1$ and $\alpha = 10$ dB/km and beamwidth=10° and different values of the SINR threshold β . The probabilities were computed analytically from (3.18) and also from (3.20) where the second term in (3.20) was evaluated by the use of Monte-Carlo simulations and the first term from (3.21) with $R_{th} = 40R_0$ and $\Delta P_{dB} = P_{0,dB} - 30$ dB.

We observe that when the desired SINR β increases beyond about 15 dB, the probability of collision approaches 10%. For the parameters corresponding to the reference link budget described in Section 3.4, the collision probabilities are small. Furthermore, in this highly directional regime, the probability of collision under

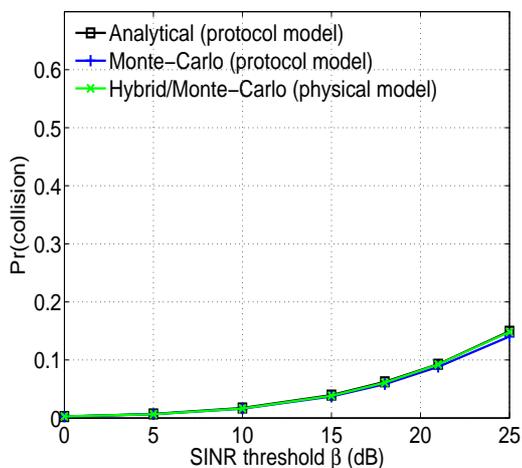


Figure 3.6: Flat-top antenna.

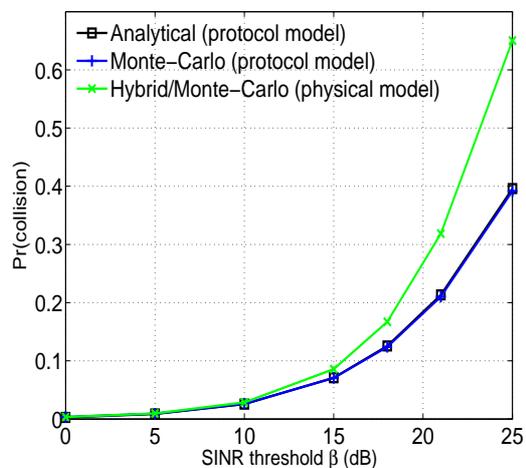


Figure 3.7: Linear array.

the physical model does not appreciably differ from the protocol model. This indicates that we may not need to worry about far-away interferers or the details of antenna beam patterns (using the notion of the equivalent flat-top beamwidth). We conclude that the MAC designer can use the following **pseudo-wired abstraction**: as a starting point: (1) *Half-duplex constraint*. Each node can either transmit or receive at any given time but not both. (2) *No interference*. Transmissions between two distinct pair of nodes are unlikely to interfere with each other, and can be largely ignored in MAC design.

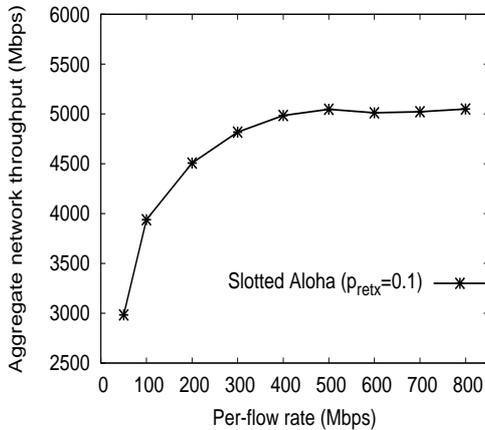


Figure 3.8: Aggregate network throughput versus per-flow rate.

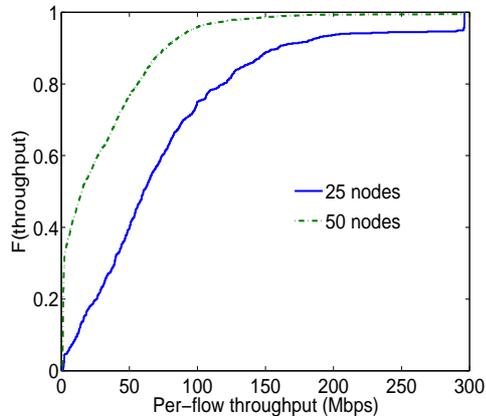


Figure 3.9: Per-flow throughput empirical CDF versus node density.

3.6 Simulations

To verify the pseudowired abstraction, we simulate a naive slotted Aloha protocol (prior studies on slotted Aloha with directional communication include [28, 75, 76], but our goal is to examine the relative effects of interference and deafness on performance). We note that far better performance can be obtained using more sophisticated MAC designs presented in the next section.

Simulation set-up: We consider random network topologies with 25 or 50 nodes spread over a 500m x 500m flat terrain. Every node initiates one constant bit rate (CBR) flow to each of its neighbors. Whenever a node has a new packet to transmit, it beamforms towards the direction of the intended receiver and transmits the packet in the next slot. If the node does not receive an ACK, it attempts to retransmit the packet with a probability p_{retx} over the next slots. A node returns

to the unbacklogged state after every successful packet transmission. We consider a sectorized antenna design with each sector covered by an array of high-gain horn elements (e.g., a linear array of four horn elements of directivity 18dBi each achieves the total directivity of 24dBi, with the horizontal sector-span of 20 degrees.) We use the QualNet Network Simulator [61], modifying the QualNet PHY and Antenna modules to model propagation in the mm-wave band and our link budget design.

Fig. 3.8 shows (for 25 nodes) the aggregate network throughput versus network load, with $p_{retx} = 0.1$. The throughput is significantly higher than with omnidirectional Aloha. Fig. 3.9 plots the empirical cumulative distribution function of the per-flow throughput for slotted Aloha for an input per-flow rate of 300Mbps for 25 and 50 node random topologies over a fixed terrain. Clearly, naive Aloha does not achieve anywhere near fair resource allocation among flows. The average throughput per flow decreases with node density: having more neighbors makes transmit-receive coordination more difficult, and increases interference. However, packet losses due to failed coordination are an order of magnitude higher than those due to interference: the fractions of failed receptions relative to the total received packets because of interference and failed coordination are 2.2% and 35.7%, respectively, for 25 nodes, and 5.6% and 47.2% for 50 nodes.

Therefore, for the directivities typical of mm wave nodes with compact form factors, the pseudo-wired model is indeed appropriate, which motivates a radically different approach to MAC design. Rather than focusing on interference management as in conventional MAC design, we must now devise scheduling mechanisms that address deafness. Based on these design principles, we now describe our Memory-guided Directional MAC (MDMAC) protocol for 60 GHz mesh networks.

3.7 Memory-guided Directional MAC

By virtue of deafness, the only feedback immediately available to a node is regarding its own transmitted and received packets. The novelty of our protocol lies in using memory and learning to arrive at approximate TDM schedules, based on this minimal information. Given the reduced spatial interference, there is no attempt by the nodes to account for the effect of their own transmissions on their neighbors. However, mechanisms for adapting the TDM schedules are built in; this allows reaction to interference (in the few cases where it does occur) and changes in the traffic pattern, and avoids locking into grossly unfair schedules. We begin our exposition with a naïve approach to using memory, and successively introduce the refinements that lead to our final design, in order to clarify the reasoning behind our design choices.

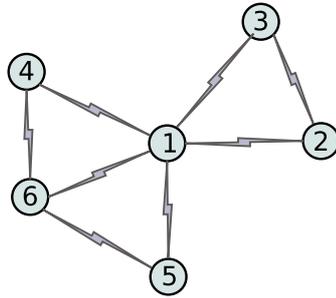


Figure 3.10: A six node network toy example.

3.7.1 Use of Memory for Implicit Coordination

Suppose that node A wishes to transmit to a neighbor node B, and randomly picks one of the free slots in a frame to do so. If the transmission is successful, node B responds with an ACK, and both nodes mark the slot as designated for communication from A to B in future frames, thus creating an *implicit reservation*. Node B uses the receive beamforming weights learnt during the first transmission in future slots corresponding to this reservation, and uses reciprocity to transmit beamform its ACKs to node A. The A→B reservation persists over multiple frames until (1) A has no more packets to send to B, or (2) A’s transmission fails repeatedly over multiple frames, or (3) A or B explicitly terminates the reservation. If A’s initial transmission to B fails (i.e., A does not receive an ACK from B), it flags the slot as “blocked” for future transmission attempts to B.

Problems with the naïve approach: Clearly, if all nodes follow the preceding procedure, the network settles into a TDM schedule without explicit coordination.

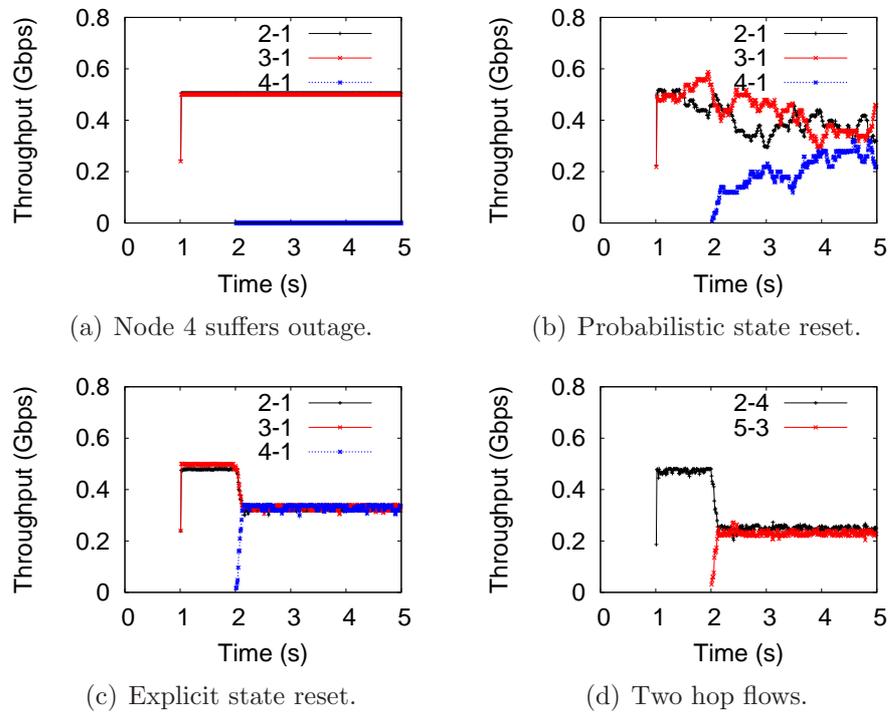


Figure 3.11: Instantaneous throughputs

However, this naïve approach can have a number of problems. Nodes which start late can get locked out in a saturated network, and deadlock conditions can arise due to, for example, nodes A and B trying to send to each other in the same slot. As a toy example, consider the simple six-node network in Fig. 3.10, where nodes 2 and 3 each have a high data rate flow to node 1 starting at 1s. Node 4 attempts to obtain bandwidth starting at 2s. Assume that all the flows have a base rate enough to saturate the 1 Gbps link. We simulate this scenario over QualNet, using simulation parameters described in detail later, in Section 3.9. Fig. 3.11(a)

shows the instantaneous flow throughputs. While the throughputs for the flows $2 \rightarrow 1$ and $3 \rightarrow 1$ settle to half the link capacity, flow $4 \rightarrow 1$ gets completely locked out because we have not yet put in mechanisms for perturbing the TDM schedules to accommodate new flows.

Probabilistic State Reset: We now introduce a decentralized, probabilistic mechanism that produces enough churn to allow rearrangements in the TDM schedules. We limit the slot state lifetimes in the following manner: over each frame, there is a nonzero probability that a slot state (e.g., transmit to, receive from, or blocked for transmit attempts to, a given neighbor) is reset, making the slot available for future frames. Thus, each node gives up some of its committed transmit or receive slots and also forgets the blocked slot information after a random number of slots. Due to the randomization in this mechanism, different slots become available at different times, thereby offering all the neighbor nodes a chance to grab these slots. The reset probabilities can be state-dependent, and can be tuned to respond to the expected traffic dynamics.

For our simple example of Fig. 3.10, we find that assigning random lifetimes to slot states (e.g., with average lifetime of 100 frames) does prevent node 4 from being completely locked out. However, as shown in Fig. 3.11(b), the throughput achieved by node 4 over the initial three seconds is still much less than that achieved by nodes 2 and 3. This motivates the need for devising more effective

ways to perturb the TDM schedules so as to quickly respond to traffic demands, and to maintain some notion of fairness. Before we describe further refinements towards this end, however, we present a simple model that offers analytical insights into the average throughput performance.

3.7.2 Approximate Protocol Modeling

For our analytical model, we assume that each node has a pseudowired link to each neighbor. In addition, we make the following simplifying approximations that decouple the complex interactions among neighbor nodes or links, and allow us to capture the essential tradeoffs in a compact manner:

- We focus attention on protocol dynamics at a single “typical” node, as an approximate representation of interior nodes in a large network. We also assume that all time slots in a frame evolve independently, and develop a Markov model for the state of a given slot over multiple frames.
- Each node maintains state information on each of its outgoing and incoming links for each slot. Outgoing links can be in one of the following states: “Transmit” (T), “Idle” (I) and “Blocked” (B), where respectively the link is actively transmitting, idle or blocked in the corresponding time slot. The MDMAC protocol allows only links in the “Idle” state in a given slot to contend for a reservation. For incoming links, this simplifies to only two states: “Receive” and “Idle”. In addition, we

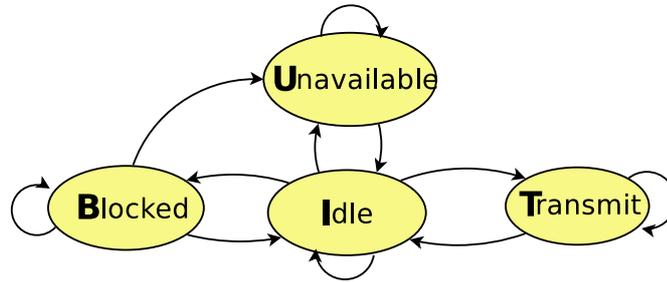


Figure 3.12: MDMAC model: state diagram for an outgoing link.

introduce the “Unavailable” (U) state, indicating that the link cannot contend for that slot because some other link for that node is in “Transmit” or “Receive” state.

- We approximate the state of the links of a node as independent of the state of the links for all other nodes. Clearly, this is not strictly true: a “Transmit” state at node A sending to node B automatically implies that node B is in “Receive” state for that slot. However, this decoupling approximation provides a convenient scenario for analysis.
- The schedule activated on each slot is chosen randomly and independently from the other slots. This is reasonable for a saturated network where all nodes always have traffic to transmit on all of their links.
- We approximate the states of different links for a given node (in a given time slot) as independent.

The state diagram for an outgoing link under this model is illustrated in Fig. 3.12. We denote the steady state probabilities of state s as P_s , where $s \in \{T, I, B, U\}$, and the transition probability from s_1 to s_2 as $P_{s_1 s_2}$. Let N denote the

number of neighbors for a “typical” node. We now introduce a new tunable protocol parameter, corresponding to *nonpersistent contention*: the *listening probability* p_l which is the probability that a node decides not to contend for transmission on any of its outgoing links. Thus, the probability that a node with all links idle chooses to transmit over a given link (each link is chosen with equal probability) is p_{tx}/N , where $p_{tx} = 1 - p_l$. The steady-state probabilities for any state s satisfy:

$$P_s = \sum_{s'} P_{s's} P_{s'} \quad (3.22)$$

Assuming that the neighbors of the typical nodes are themselves typical, the transition probabilities for the Markov model in Fig. 3.12 must satisfy certain consistency conditions. Before developing an iterative procedure for computing these in general, we first consider a simple two node network.

Modeling a Two-Node Network: For the outgoing link from a node, we observe the following: the Unavailable state means that the incoming link from the neighbor is active, the Idle state means that the incoming link is Blocked or Idle, and the Blocked state means that the incoming link is Blocked or Idle. We now compute the transition probabilities for one of the links. The probabilistic state resets are modeled as follows: T_{slot} and T_{block} are the average transmit/receive slot and blocked slot lifetimes. Therefore,

$$P_{TI} = \frac{1}{T_{slot}}, \quad P_{UI} = \frac{1}{T_{slot}}, \quad \text{and} \quad P_{BI} = \frac{1}{T_{block}} \quad (3.23)$$

To compute P_{IT} , note that this state transition requires that the node choose to contend in the slot rather than listen (probability p_{tx}) AND either the receiving node is in Blocked state, OR the receiving node is in Idle state AND chooses to listen (probability p_l). Thus we have:

$$P_{IT} = p_{tx} \left(\frac{P_I}{P_I + P_B} p_l + \frac{P_B}{P_I + P_B} \right) \quad (3.24)$$

where we used the independence approximation to infer that the conditional probabilities of the incoming link being in the Blocked or Idle states given the reference link's Idle state are proportional to the respective steady-state probabilities. Using similar reasoning, we can evaluate the other probabilities as:

$$P_{IU} = \frac{p_{tx} p_l P_I}{P_I + P_B}, P_{BU} = \frac{p_{tx} P_I}{P_I + P_B}, P_{IB} = \frac{p_{tx}^2 P_I}{P_I + P_B}. \quad (3.25)$$

Procedure 1 defines an iterative algorithm to compute the steady-state probabilities. Consistency demands that the steady-state probabilities of the Transmit and Unavailable states should be equal for the two node network because of symmetry, and indeed we find that this is always the case. We now compare the analytical and simulation results for the expected link utilization for successful transmissions for each node. Observe that the total medium utilization in this case would actually be the sum of the transmit link utilization for the two nodes sharing the link capacity. The steady-state probabilities calculated from the Markov chain model are: $P_T=P_U = 0.489$, $P_I = 0.015$ and $P_B = 0.007$. QualNet simulations of

Algorithm 1 State probabilities computation

- 1: Initialize $P_I = 1$, $P_T = P_U = P_B = 0$.
 - 2: Use the current values of P_s and (3.23), (3.24), (3.25) to compute the transition probabilities $P_{ss'}$.
 - 3: Use the values for $P_{ss'}$ obtained in Step 2 and (3.22) along with the normalization condition $\sum_s P_s \equiv 1$ to solve for the state probabilities and update the values of P_s .
 - 4: Return to Step 2 until convergence.
-

the protocol for this setting yields the fraction of the successful transmit and receive state slots as 0.492 each, which demonstrates a close match. The other state fractions are Blocked: 0.013 and Idle: 0.002 - the differences correspond to additional refinements embedded into the actual MDMAC QualNet protocol model (see Section 3.7.4).

The preceding reasoning can be extended to obtain the state probabilities $\{P_s\}$ of the “typical” link for an arbitrary network, i.e., for nodes with N bi-directional links for arbitrary N . The algebra is a bit more involved. For instance, unlike for the two-node network, a link being in Unavailable state does not automatically mean that the corresponding incoming link is active; any of the other $2N - 1$ links from/to the same node can be active. We illustrate the method by deriving an expression for P_{IT} , the probability of transitioning from Idle to Transmit state on

the reference link. This requires that the node chooses to contend for transmission (probability p_{tx}) rather than to listen, and that out of the subset of the N outgoing links that are in Idle rather than Blocked state, the reference link is chosen. We already know that the reference link is in Idle state. The probability of any other link being in Idle state is $P_I/(P_I + P_B)$. Thus, the probability of the reference link being chosen to contend is given by

$$p_c = p_{tx} \sum_{k=0}^{N-1} \binom{N-1}{k} \left(\frac{P_B}{P_I + P_B} \right)^m \left(\frac{P_I}{P_I + P_B} \right)^k \frac{1}{k+1},$$

where $m = N - 1 - k$. In addition, for the contention attempt to be successful, the corresponding receiving link on the neighboring node, must be in the Idle or Blocked states, and the neighboring node must choose to listen rather than contend to transmit itself on one of its Idle outgoing links (if there are any Idle links, or else all its outgoing links must be in Blocked state). The corresponding probability (p_r) is

$p_r = p_l(P_I + P_B(1 - \left(\frac{P_B}{P_I + P_B}\right)^{N-1})) + P_B \left(\frac{P_B}{P_I + P_B}\right)^{N-1}$. Furthermore, the receiving node must choose the reference link out of all of its own other $N - 1$ neighboring nodes (also “typical” nodes) who also happen to be Idle and contending for this slot (probability $P_I p_c$). Using this we finally have

$$P_{IT} = p_c p_r \sum_{k=0}^{N-1} \binom{N-1}{k} \frac{(P_I p_c)^k}{k+1} (1 - P_I p_c)^m, \quad (3.26)$$

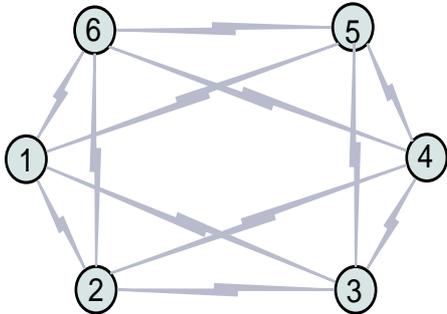


Figure 3.13: An example network with 4 neighbors per node.

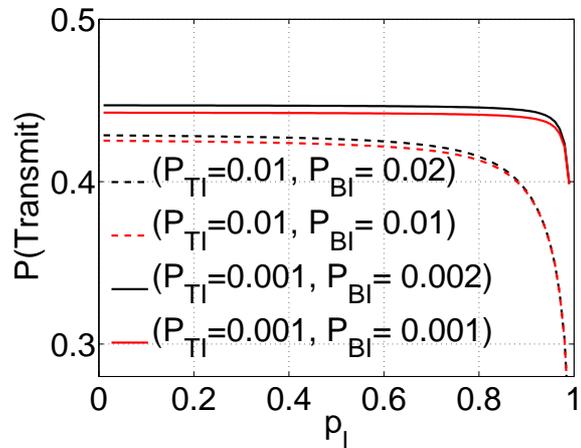


Figure 3.14: Medium utilization by node transmissions.

where $m = N - 1 - k$. The expressions for other transition probabilities can be computed using similar reasoning: we omit the details here for lack of space. Procedure 1 can then be used to compute the steady-state probabilities $\{P_s\}$. Note that consistency requires that $P_U = (2N - 1)P_T$.

To illustrate the results of the analysis, let us first consider a six-node topology in Fig. 3.13 where every node has four neighbors with link-saturating flows in each direction. The steady-state medium utilization for successful transmissions from each node is obtained as 0.426. Packet-level simulation of the same scenario yields 0.43. We now employ our analytical model to get insight into the effect of the following parameters on the performance of MDMAC (1) P_{TI} , P_{UI} , and P_{BI} ; and (2) $p_l = 1 - p_{tx}$, under saturated traffic conditions and a given network node density. We consider the probabilities P_{TI} and P_{UI} such that the average transmit/receive

state lifetime T_{slot} is in the range of 100 to 1000 frames, which correspond to the P_{TI} , P_{UI} values of 0.01 to 0.001. We choose these values based on the criterion that the schedules should be maintained for a time that is one or two orders of magnitude larger than the worst-case packet round trip times for nodes which are multiple hops away from a gateway. In order to understand the effect of tuning the state lifetimes, for each value of P_{TI} , we consider the following Blocked state reset probabilities: $P_{BI} = P_{TI}$ or $P_{BI} = P_{TI}/2$. For the symmetric six-node topology, Fig. 3.14 shows the expected total medium utilization by a node's transmissions to all its neighbors. We only show a sample of our results to highlight the insights. We find that $P_{TI} = 0.001$, and $P_{TI} = 0.002$ yields the best performance among the parameter choices that we have considered. Also, the medium utilization is relatively insensitive to the value of the listening probability p_l (in the range 0.1-0.7). We observe similar trends for medium utilization over a large range of neighbor densities, and therefore set $P_{TI} = P_{UI} = 0.001$ and $P_{BI} = 0.002$ for our evaluation (Section 3.9). These relatively large state lifetimes result in less throughput loss due to churn, while (as we shall see) providing enough possibility for rearrangement of schedules to ensure fairness. Although our analytical guidance for choice of P_{TI} , P_{UI} and P_{BI} is for saturated traffic, our simulations indicate that these values are effective for unsaturated multihop mesh traffic as well.

3.7.3 Fairness: Explicit State Reset (ESR)

In addition to the preceding mechanisms of probabilistic state reset and non-persistent contention, we introduce the following strategy to improve the fairness of link bandwidth allocation among neighbors: as long as the fraction of committed transmit or receive slots at a node is less than a threshold (say, $T_{ESR} = 90\%$), its medium access continues as before. However, when this fraction exceeds T_{ESR} , the node successively picks the neighbor that holds the highest share of the committed transmit or receive slots, and resets the state (i.e., considers the slot as free) for a randomly picked slot among the committed slots until the fraction drops below the threshold. Transmit and receive commitments are treated separately while deciding which slot to free up: for many mesh traffic scenarios, the allocation can otherwise be very asymmetric in terms of the transmit and receive slot allocations to a neighbor, which can lead to gross unfairness. The ESR mechanism ensures that when the bandwidth demand at a node starts to approach the total link capacity, the node switches to an “alert” mode to ensure that the bandwidth allocations are not grossly unfair. The node perturbs the TDM schedule by freeing some slots from neighbors that consume the greatest share of the link bandwidth, in order to free up bandwidth for other neighbors. There is no need for this mechanism when the total demand is considerably less than the link bandwidth, but we do need enough “slack” in the schedules such that the probability

that a node quickly grabs a free slot is reasonably high. Based on this tradeoff, we have found from our simulations that a 90% threshold works well in practice, in helping competing neighbors obtain a near-fair share of transmit and receive opportunities under heavy traffic.

Fig. 3.11(c) shows that the instantaneous flow throughputs with the ESR strategy employed for our toy example (Fig. 3.10) correspond to a near-fair bandwidth allocation, when nodes 2, 3 and 4 are contending for bandwidth to transmit to node 1. Fig. 3.11(d) shows the effect of ESR for multihop flows over our toy network: flows 2→4 and 5→3 start at 1s and 2s, respectively, and are seen to settle into a fair schedule. More sophisticated rules could be applied to decide which slots to free up in order to enforce alternative bandwidth allocation policies (e.g., proportional fairness), but this is beyond our current scope.

3.7.4 MDMAC Protocol Design Details

Slot allocation tables: Each node maintains a slot allocation table that contains its current transmit/receive state information. The state information for a particular slot is updated either at the end of the slot, or at the beginning of a new frame: for example, after a successful transmission on a previously idle slot, the slot's state changes from Idle to Transmit, and the associated neighbor information is also stored. For N neighbors, and k slots per frame, the additional memory required to

store the state information is $O(Nk)$ units, which is dominated by the worst-case size constraint on the blocked slot list. Thus, the memory and information processing requirements for the protocol do not represent an implementation bottleneck.

Blocked slots: If a node's transmit attempt fails, the slot is blocked for a geometric number of frames, governed by the probability P_{BI} . A blocked slot is not used for future transmit attempts to that neighbor, *except* when there are no unblocked free slots. In that case, blocked slots are picked with a low probability from the pool of free blocked slots. This mechanism can be viewed as a smart per-node directional backoff.

Persistence on packet loss on an existing schedule: If a packet transmission fails in the Transmit state, the node changes the state to Tx-Unsure. The node will continue to transmit over the slot, switching back to the Transmit state on a successful packet transmission over the next frames, or freeing up the slot if the packet transmission continues to fail. The state transitions for the receive side are similar: the corresponding states are Receive and Rx-Unsure. The states Tx-Unsure and Rx-Unsure control the level of persistence in the face of failed transmissions *after* the schedule is set to Transmit/Receive: they help to ensure that the schedules are relatively stable even under occasional packet loss. Nodes explicitly communicate to the corresponding neighbor whenever a slot state (Transmit/Receive) is reset to Idle, either because of the probabilistic state reset or the ESR mechanism: this

is done via a state-reset bit in the data packet/ACK message headers, or via an ESR message if the transmit packet queue is empty. This communication prevents slot wastage on state reset by avoiding the neighbor's transition to the persistence states described above.

Per-neighbor packet queue backlog-based slot contention: Each node maintains a per-neighbor output queue, and decides to contend for more transmit slots to a neighbor based on the current queue backlog. For example, if the queue backlog to a neighbor exceeds a threshold T_Q , the node decides to contend for a new slot. Any slot in the Transmit state is freed if the corresponding neighbor queue is empty.

3.8 Medium Utilization with MDMAC

Before delving into detailed simulations that incorporate physical interference, we first examine the properties of the TDM schedules resulting from MDMAC, via the pseudowired model. Consider a directed graph $G = (V, E)$ where V is the set of nodes and the set of edges E represents the directed wireless links between directly communicating neighbor nodes. We define a *feasible schedule* $F \in E$ as a set of links that can be simultaneously activated, i.e., that do not violate the interference constraints (for the pseudowired model this is the half-duplex constraint).

We then define a maximal schedule $M \in E$ as a feasible schedule that has the additional property that adding any additional link to the schedule will result in an infeasible schedule. Assume that every node has a link-saturating flow to each of its neighbors. A TDM algorithm is a procedure that maps the set of available slots to the set of feasible schedules. Clearly, any optimal TDM schedule requires that a maximal schedule be used on every slot. While the MDMAC protocol does not guarantee the use of a maximal schedule over every slot, our numerical results show that it does achieve close to optimal schedules with high probability.

Worst-case missed transmit opportunities: We first consider a worst-case estimate of the “missed” transmit opportunities under MDMAC. Specifically, we compare MDMAC against a genie-based procedure in which the genie first looks at the set of successful links activated by MDMAC, and then uses the complete network topology information to schedule transmissions on the largest-cardinality maximal matching of links that is a superset of the links activated by MDMAC. To compute this, we consider a reduced graph R_i of links that can be simultaneously activated with the set of links in the MDMAC schedule. Since MDMAC’s schedule over each slot is typically close to a maximal matching, R_i is usually a very small graph, for which the cardinality of the maximum matching can be quickly obtained using a polynomial time maximum matching algorithm as in [77]. With K links in the MDMAC schedule and L additional links found by the genie, the

fraction of missed transmit opportunities is $K/(K+L)$. If there was no successful transmission over a slot, we randomly pick a maximal matching for comparison, and set the fraction of missed transmit opportunities to one. For our six node toy network (Fig. 3.10), we find that on average, MDMAC misses only about 5% of transmit opportunities as compared with the genie-based approach.

Comparison with Greedy Maximal Scheduling (GMS): We also compare MDMAC with a centralized GMS algorithm similar to [78], with the link weights proportional to the fraction of past slots the link was inactive. The latter is a more meaningful weight measure for saturated traffic than the link queue lengths used in [78] as finite-size packet queues would overflow under saturated traffic. The basic idea behind GMS is to incrementally obtain a maximal schedule in the following manner: pick a link l with the maximum weight (breaking ties arbitrarily) among the set of candidate links C (initially set E), remove l and the links that cannot be simultaneously activated with l from the set C , continuing until C is empty. Note that both centralized and distributed versions of GMS have significantly higher control overhead than MDMAC (e.g., see [78]). Our objective here is simply to show that MDMAC provides competitive medium utilization despite its simplicity: for our toy topology (Fig. 3.10), the expected number of links used per slot for the GMS algorithm turns out to be 2.4, whereas on average, MDMAC uses 2.43 links per slot.

The preceding performance metrics are calculated based on the pseudowired model, which completely ignores signal interference (and hence overestimates performance). We next report on packet-level simulations incorporating a detailed 60 GHz physical layer model that captures topology-dependent (albeit low) interference losses and random packet loss due to noise.

3.9 Simulation Evaluation

We compare the performance of MDMAC against a directional Slotted ALOHA (DSA) protocol that benefits from the reduced interference by exploiting the highly directional 60 GHz links (see Section 3.6). DSA serves as a good baseline to understand the potential performance advantages of using memory for medium access coordination. We do not undertake a comparison of MDMAC with the other lower frequency directional MAC protocols, since these protocols are either not applicable because of their underlying assumptions or not designed for the highly directional 60 GHz networks that we consider. For example, most WiFi mesh networking protocols primarily focus on interference management and incur significant control overhead (e.g., RTS-CTS type message exchange, out-of-band tones, directional message broadcasts in each sector, etc.) to avoid interference (see Section 3.1).

Parameter	Value
PHY raw data rate, uncoded QPSK link (R)	2 Gbps
PHY overhead (signal acquisition/beamforming) (T_{PHY})	$1\mu s$
Payload + header transmit duration	$(1000 + 56) * 8/R$
SIFS interval	500ns
Total packet transmission+ACK duration	$7.69\mu s$
MDMAC slot duration, frame duration (50 slots)	$8\mu s, 400\mu s$
$P_{TI} = P_{UI}$, and P_{BI}	0.001, 0.002
Queue backlog to trigger a new slot contention	≥ 6 packets
Prob(a blocked slot picked when no other free slots available)	0.02

Table 3.1: Simulation parameters

In addition to the utilization-related performance metrics described in Section 3.8, we define a weighted *MAC Fairness Index (MFI)* metric in order to accurately capture the MAC level resource allocation fairness in a saturated network:

$$MFI = \frac{(\sum_{l=1}^m y_l/w_l)^2}{m \sum_{l=1}^m y_l^2/w_l^2} = \frac{\mu^2(y_l/w_l)}{\mu^2(y_l/w_l) + \sigma^2(y_l/w_l)}, l \in E$$

where $w_l = 1/\max(cs(l), cd(l))$, $cs(l)$ and $cd(l)$ are the number of neighbors of link l 's transmit and receive nodes, respectively; m is the total number of links, and μ and σ are the mean and the standard deviation of the weighted link flow throughputs y_l/w_l over all links l . Note that $0 \leq MFI \leq 1$, and for $w_i = 1 \forall i$, MFI equals the Jain's Fairness Index [79].

Simulation Model: We evaluate the performance of MDMAC and DSA via simulations over the QualNet Simulator. We have modified the QualNet physical layer and Antenna modules to model the 60 GHz propagation characteristics and capture our link budget (see Section 3.3). The transmit/receive antenna patterns

are obtained from a simple abstraction of the antenna models discussed in Section 3.1. We divide the 360° azimuthal coverage area into 18 sectors. Each sector is covered by a linear array of four highly directive antenna elements (with element pattern approximated by a horn element pattern of directivity 18 dBi). At any instant, the steered antenna patterns and other link budget parameters are used to calculate the signal to interference and noise ratio (SINR) at a node, which then yields the bit error rate (BER) based on the modulation: the packet loss probability evaluated from the BER and the packet size determines successful or failed packet reception. We note that our physical and MAC layer simulation parameters listed in Table 3.1 are representative numbers used for illustration; our inferences about MDMAC performance do not depend on these specific parameter choices.

We consider random mesh network topologies of 25 and 50 nodes spread over a 500m x 500m flat terrain. The 50 node topologies are considered to stress-test our protocol design for high-density deployments where interference becomes more significant, and coordination is more challenging. For the mesh topologies with gateway nodes, we run a clustering algorithm over the random topologies to assign gateway nodes in such a way that no mesh node is more than three hops away from its gateway. We simulate the network discovery procedure during the network initialization phase from 0 to 1s. The start times of all the flows are

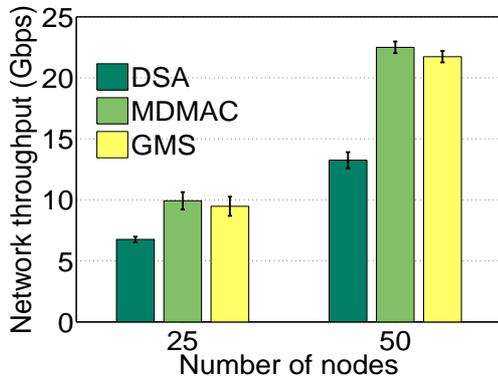


Figure 3.15: Aggregate throughput.

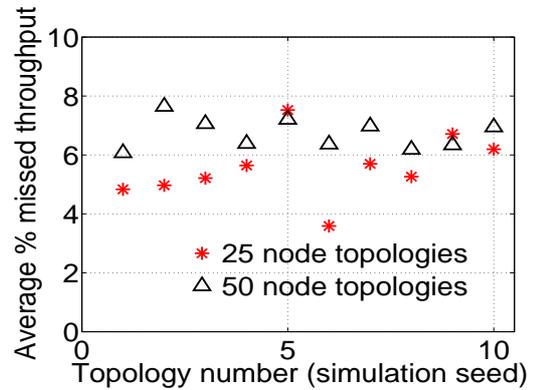


Figure 3.16: Missed transmit opportunities.

	25 node topologies	50 node topologies
GMS	0.93	0.93
MDMAC	0.91	0.88
DSA	0.39	0.20

Table 3.2: MAC Fairness Index (MFI)

randomly chosen between 1 to 2s. Our simulation results are obtained over 10 simulation runs with a different network topology and traffic for each seed.

Saturated Network Traffic Model: Fig. 3.15 compares the aggregate network throughput achieved by GMS, MDMAC and DSA. The aggregate network throughput for MDMAC is comparable to the centralized GMS, and much higher than DSA. Table 3.2 compares the MAC Fairness Index (MFI) averaged over the different simulation runs. MDMAC’s MFI is slightly lower than the MFI for GMS; however, both values are reasonably close to the optimal value of 1. For DSA, the lack of coordination leads to unacceptable link outages for a large frac-

tion of neighbors (20% and 40% on average for the 25 and 50 node topologies, respectively), which translates to DSA's low MFI.

Fig. 3.16 shows MDMAC's average worst-case missed transmit opportunities: each corresponding to a different simulation experiment, calculated using the procedure described in Section 3.8. We note that on average, the TDM schedules generated by MDMAC are within 6% and 7% of the corresponding largest cardinality maximal matchings on the network graphs for the 25 and 50 node topologies, respectively. The low missed transmit opportunities demonstrate the efficiency of the distributed TDM schedules that emerge from MDMAC.

Multihop Mesh Traffic Model: This model captures the traffic asymmetries expected in multihop mesh networks. We assign a UDP-CBR flow of rate 200 Mbps from each mesh node to its gateway or/and from a gateway to each node in its cluster, with a probability 0.5. Fig. 3.17 shows the aggregate network throughput for MDMAC and DSA. MDMAC's implicit transmission coordination results in a significantly higher aggregate throughput than DSA (35% and 52% higher for the 25 and 50 node topologies, respectively). The higher fractional gain for MDMAC over the 50 node topologies demonstrates that MDMAC is effective despite the increased contention and interference resulting from the high node density.

We now look at the end-to-end delay and delay jitter for the received packets to obtain insight into the QoS performance. Fig. 3.18 plots the average packet delay

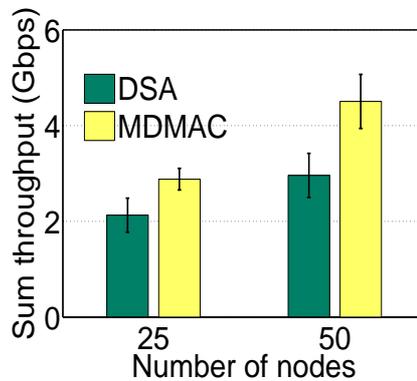


Figure 3.17: Throughput.

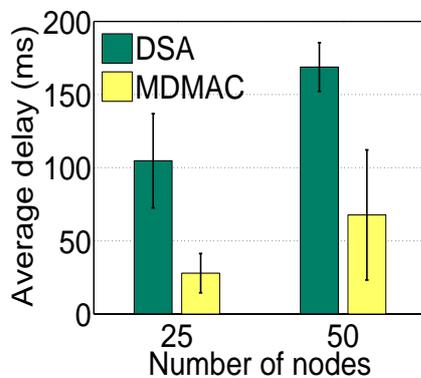


Figure 3.18: Total delay.

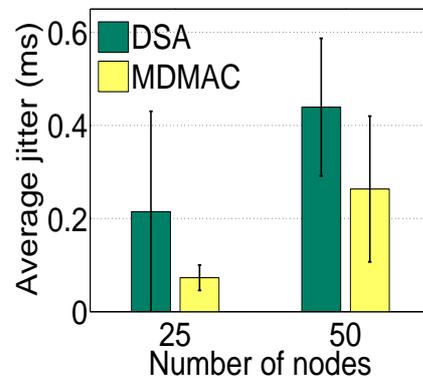


Figure 3.19: Delay jitter.

for MDMAC and DSA. Observe that the times taken by the nodes to find slots in MDMAC's schedules over multiple hops turn out to be much lower than the typical Internet traffic delay requirements which are on the order of 100s of milliseconds. MDMAC's TDM-like performance leads to much lower jitter, as seen in Fig. 3.19.

To illustrate MDMAC's ability to adapt to changes in the traffic pattern, we consider an example with short-lived, high datarate flows, over a portion of a 25 node random topology around a gateway node 9 (see Fig. 3.20). Fig. 3.21 shows

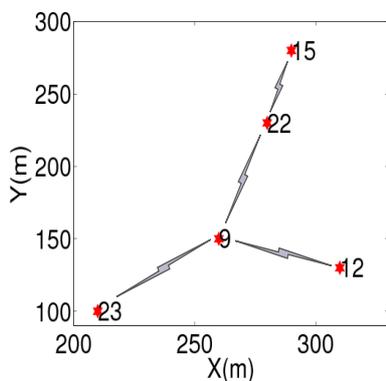


Figure 3.20: Portion of a 25 node mesh topology.

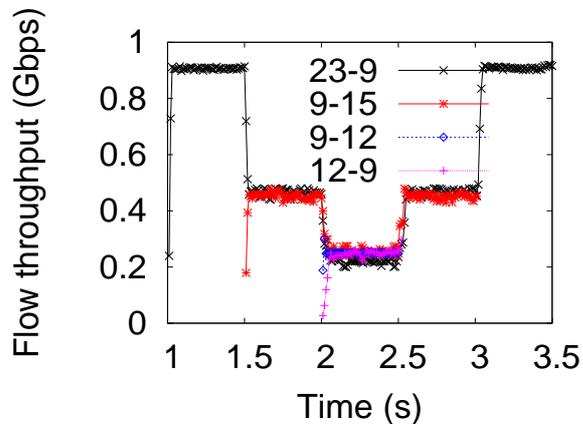


Figure 3.21: Instantaneous flow throughputs.

the instantaneous flow throughputs: flow 23→9 starts at 1s with a throughput of 900Mbps. At 1.5s, flow 9→15 of lifetime 1.5s starts: the schedules quickly adapt such that the two flows share the link bandwidth equally. At 2s, a bidirectional flow (9→12, 12→9) of lifetime 500ms is initiated: MDMAC’s explicit state reset mechanism adapts to the changed traffic and readjusts bandwidth allocations for a fair sharing of the total bandwidth. The flow throughputs quickly recalibrate after a flow ends. The results illustrate MDMAC’s ability to quickly adapt and maintain link-level fairness.

3.10 Summary

We have pointed out how the unique physical layer characteristics of mm wave links impact MAC design for 60 GHz outdoor mesh networks. Our interference

analysis framework enables a quantitative evaluation of when we can model highly directive links as pseudo-wired. For the directivities typical of mm wave nodes with compact form factors, the pseudo-wired model is indeed appropriate, which motivates a radically different approach to MAC design. Rather than focusing on interference management as in conventional MAC design, we must now devise scheduling mechanisms that address deafness. Armed with this insight, we then showed that 60 GHz mesh networks can operate at high utilization using a completely distributed MAC that exploits the interference reduction, while dealing with the deafness that arises due to the use of highly directional, electronically steerable, beams. The key idea is to obtain TDM-like schedules by using memory about which slots work and which do not, while providing enough random “churn” to allow for quick schedule adaptation. Finally, the promising results in this chapter indicate the potential gains from memory and learning-based distributed coordination in more general settings, including omnidirectional networks.

Chapter 4

Sticky CSMA/CA

4.1 Introduction

As we discussed in Chapter 1, wireless mesh networks offer an attractive solution for providing voice, multimedia, and best effort data services in areas where deployment of a wired infrastructure is not viable or is economically unattractive. However, a key drawback of current wireless mesh technology based on IEEE 802.11 standards is that it is not possible to provide bandwidth and delay guarantees to real-time applications such as voice and video [80]. The primary reason is the variability in delay and loss performance that results from the requirement that each packet must contend afresh for the medium. These drawbacks of IEEE 802.11 are only partly alleviated by its QoS extension IEEE 802.11e. Given that medium access control is the bottleneck for providing QoS, an alternative approach is to employ Time Division Multiplexing (TDM) to obtain a collision-free

transmission schedule. However, this requires fine-grained network-wide synchronization and scheduling, which are difficult to implement. In this chapter, we present Sticky CSMA/CA¹, a medium access control mechanism that provides TDM-like performance to real-time flows without requiring explicit synchronization.

Our starting point is the observation that the traffic generated by real-time flows exhibits a significant regularity. Can this regularity be exploited for more effective medium access control? After all, contention between nodes occurs because a node cannot predict the actions of its neighbors. If the transmission schedule for a node follows a regular pattern, its neighbors should be able to learn this pattern using the CSMA mechanism, and thus avoid contention. While these ideas are potentially applicable to a broad class of real-time flows, in this chapter, we focus specifically on illustrating how they enable efficient support of Voice over IP (VoIP) over wireless mesh networks. The success of VoIP over the wireline Internet leads us to expect that it will be a key application driving the wider acceptance and commercial success of wireless mesh networks. A broad range of wireless VoIP solutions are commercially available and are being deployed in enterprises and campuses [82, 83]. However, these products typically optimize VoIP

¹The term *sticky routing* has been used for routing protocols in telephone networks [81] that grab and hold resources while feasible. Sticky CSMA/CA also grabs and holds resources, but for channel access: the context and protocol details are completely different.

performance over a single wireless hop from a tetherless node to an Access Point connected to the wired network. Our goal is to support VoIP in wireless mesh networks, thus significantly increasing the flexibility of deployment. The challenge is to provide an adequate level of QoS, in terms of assured bandwidth, and bounded delay and delay jitter.

Sticky CSMA/CA is based on the assumption that all the real-time flows in the network are either naturally periodic (e.g., VoIP) or have been shaped by the higher layers as periodic (constant bit rate) streams with the same period. A node that has to transmit packets from a real-time flow monitors the medium using the standard CSMA mechanism. When it detects an opportunity to transmit, it attempts packet transmission, and on being successful, *sticks* to a periodic schedule. Neighbors detect such periodic schedules using the CSMA mechanism and avoid interfering with them. This medium access approach leads to a TDM-like sharing of the medium among the real-time flows, with the periodicity of the application and transmission schedule being used to obtain an implicit form of synchronization. Delay-insensitive traffic has a lower priority than real-time traffic and utilizes the bandwidth left over from the real-time flows.

From the point of view of implementation, Sticky CSMA/CA requires that the nodes maintain *carrier sense tables* that record the periodic flows whose existence it can infer using the CSMA mechanism. Upon arrival of a new real-time flow, a

node looks at its carrier sense table to determine a periodic schedule over which the medium is free, and contends for and grabs the medium at the first such opportunity (i.e., it starts sending in a time interval that does not interfere with the existing real-time flows in its carrier sense table). If the flow setup is successful, the node locks onto a periodic transmission schedule. All the other nodes in the network that overhear the control message exchange update their carrier sense tables immediately, and the nodes that sense the new periodic transmission infer that a new flow has been set up. Delay-insensitive traffic has lower priority than real-time traffic and fills in the gaps remaining after the real-time flows have been accommodated. We develop new contention resolution mechanisms for delay-insensitive traffic that enhance the efficiency with which they utilize the left-over bandwidth.

Our performance evaluations focus on a mix of VoIP and data traffic, for which large gains in VoIP call carrying capacity are demonstrated relative to both IEEE 802.11b and IEEE 802.11e. We show through analysis and extensive simulations that Sticky CSMA/CA achieves more efficient medium utilization than IEEE 802.11e and can be used to provide the required QoS in terms of bandwidth, delay and delay jitter to real-time flows. The underlying reason is that this scheme obviates the need for medium contention for every frame that a node transmits, exploiting the predictable and periodic nature of the high priority traffic. This ap-

proach significantly reduces the probability of collision and the overhead due to the backoff required for collision avoidance for every frame in IEEE 802.11. We also show that, along with efficient support of significantly more real-time flows, the performance in terms of throughput achieved for the low priority delay-insensitive traffic is comparable to that of IEEE 802.11. These encouraging results indicate that even for applications that do not exhibit periodicity and constant data rates naturally (e.g., video, or even high-speed *data pipes* in a wireless backhaul), it may be useful to artificially impose periodicity and predictability to improve the efficiency of medium access. Detailed investigation of cross-layer architectures that shape variable rate real-time traffic to take advantage of Sticky CSMA/CA is beyond the scope of this chapter. However, we do provide some discussion as to how this might be done for applications such as video in Section 4.6.

The rest of this chapter is organized as follows. Section 4.2 describes related work. The design of Sticky CSMA/CA is described in Section 4.3. In Section 4.4, we present an approximate analysis to obtain insight into the performance of IEEE 802.11 and Sticky CSMA/CA. Section 4.5 presents the simulation results obtained in mesh networks with line and grid topologies. We conclude with a discussion of various design and application aspects of Sticky CSMA/CA in Section 4.6. Note that we use the terms VoIP and voice interchangeably throughout this chapter.

4.2 Related Work

In the past few years, a considerable amount of work has addressed the area of MAC layer design for providing the required QoS to voice, multimedia and data applications over wireless networks. We discuss some of the proposed approaches next.

The IEEE 802.11 standard MAC sublayer provides two different access mechanisms: the Distributed Coordination Function (DCF), and the Point Coordination Function (PCF) [84]. PCF is a centralized approach to perform bandwidth reservations, which splits time into a contention-free period (CFP) and a contention period (CP). A polling scheme controlled by a Point Coordinator that resides in the Access Point arbitrates channel access during the CFP. PCF is not suited for multihop wireless networks due to the requirement of centralized control and the Point Coordinator at the Access Point. The Distributed Coordination Function (DCF) is based on CSMA/CA for frame transmissions and a random backoff mechanism to avoid packet collisions. We describe this scheme in more detail in Section 4.3. DCF's drawback is that it does not provide service differentiation. Due to lack of both medium access priorities and QoS support for different applications, DCF is not suited for supporting real-time applications.

The IEEE 802.11e standard extends the IEEE 802.11 framework by including mechanisms for service differentiation [62]. Four access categories (AC), mapped to eight traffic priorities, are defined for IEEE 802.11e. These ACs correspond to different values of the following parameters: minimum contention window size (CW_{min}), maximum contention window size (CW_{max}), and the Arbitration Interframe Space (AIFS) (time-interval between idle transition of the medium and the start of channel access or backoff). A lower value of these parameters causes shorter backoff intervals or wait periods before medium access, thereby prioritizing traffic. IEEE 802.11e has different queues for different ACs. It also provides an optional token-like mechanism called TXOP (Transmission Opportunity) to support packet bursts, where a node that has obtained access to the medium can send multiple frames for an AC, up to a maximum time-limit, without further contention. This contention-free bursting mechanism aids in transmission of bursty traffic and improves medium utilization. One of the drawbacks of IEEE 802.11e is the lack of assured QoS for individual flows. IEEE 802.11e can only provide service differentiation among different ACs, but QoS degrades in the presence of multiple contenders in the same AC because of medium access contention with equal priority for every packet belonging to the same AC. The underlying reason for lack of support of a deterministic QoS in IEEE 802.11e and IEEE 802.11 is that every packet requires contention for the medium in these schemes, and no

attempt is made to exploit the (natural or possibly imposed) predictability of packet transmissions to minimize contention losses.

Blackburst [85] is a scheme based on modified IEEE 802.11 CSMA/CA designed to provide QoS to real-time applications. In this scheme, a node with real-time data contends for medium access by sending an energy burst (called a blackburst) for time proportional to the wait time of the node. The node with the longest burst wins the contention and transmits a frame. Blackburst does not address the hidden terminal problem and requires that the nodes have the capability to jam the medium. Another distributed MAC scheme derived from IEEE 802.11 DCF is Distributed Fair Scheduling (DFS) [86], which is based on Self-Clocked Fair Queuing (SCFQ) [87]. DFS achieves service differentiation by assigning different backoff intervals to packets that belong to different flows based on the weights assigned to the flows. This feature ensures higher throughput to flows with higher weights, and fairer overall allocation of bandwidth as compared to IEEE 802.11 DCF. DFS does not consider the delay bounds of real-time packets, nor does it address mechanisms for determining appropriate weights for different flows, both of which greatly influence the performance.

Multiple Access Collision Avoidance with Piggyback Reservations (MACA/PR) [88] is a scheme that supports bandwidth reservations for real-time traffic. In this scheme, each node maintains the real-time scheduling information of the neigh-

borning nodes in its reservation tables (RT) by overhearing the data packets or ACKs. The global reserved slot information is obtained by periodic RT exchange among neighbors. The RT exchange takes care of the hidden terminal problem. The drawbacks of MACA/PR are the piggybacking-based reservation approach and the overhead of periodic RT exchange.

A few recent proposals attempt to take advantage of TDMA schedules for reducing interference and for fair resource allocation over CSMA based networks [89, 90]. The Overlay MAC Layer (OML) [89] is an access control and scheduling scheme that partitions time into equal size slots and allocates these slots among loosely synchronized contending nodes. Though OML implements temporal fairness and reduces interference by improving the predictability of medium access among nodes, the coarse time-slot based resource allocation does not address the QoS requirements of different applications. Another interesting MAC framework proposed in the context of sensor networks is ZMAC [90]. ZMAC exploits TDMA schedules among locally synchronized neighbors to improve medium utilization and reduce packet collisions among two hop neighbors in a CSMA network operating under high contention. Under low contention, all the nodes contend for a slot but the slot-owner nodes have a higher priority over their neighbors by virtue of a shorter contention window. This scheme incurs initial schedule allocation overhead in the network deployment phase and, because it was designed

for a sensor network, it does not consider service differentiation among different applications for QoS support.

To the best of our knowledge, the proposed Sticky CSMA/CA scheme is the first approach that exploits the characteristics of real-time applications to devise medium access control strategies that reduce delay and delay jitter. Exploiting the natural periodicity of VoIP is shown (in the succeeding sections) to lead to large capacity gains relative to IEEE 802.11. This motivates future work (see the discussion in Section 4.6) in taking this approach one step further, by imposing artificial regularity at the application layer in order to obtain the enhanced QoS obtained using Sticky CSMA/CA medium access.

4.3 The Sticky CSMA/CA Protocol

In this section, we describe the Sticky CSMA/CA protocol. We briefly outline conventional CSMA/CA and introduce history considerations into CSMA/CA. We then describe the carrier sensing mechanism and discuss the operation of Sticky CSMA/CA for real-time and delay-insensitive traffic.

4.3.1 Conventional CSMA/CA

The IEEE 802.11 DCF uses CSMA/CA for data packet transmission, along with a random backoff mechanism to avoid collisions. A node that has a frame to transmit senses the transmission activity in the medium. If the medium is free for a distributed interframe space (DIFS) interval, the frame is transmitted. If the medium is sensed as busy, the node waits until the medium is free for a fixed time interval (DIFS if the last frame was received without any transmission errors, or extended interframe space (EIFS) otherwise). After the medium is sensed idle for a DIFS/EIFS interval, the node generates a backoff counter b_i , which is a pseudo-random integer with a uniform distribution in the interval $[0, CW_i]$, where $CW_i = \min(CW_{max}, 2^{i-1}(CW_{min} + 1) - 1)$ for the i^{th} transmission attempt for a frame. The node then decrements the backoff counter b_i by one for every *slot width* interval if the medium remains idle. If the medium becomes busy before b_i reaches zero, the backoff procedure is frozen and is continued after the medium is sensed idle again for a DIFS interval. When the backoff counter reaches zero, the node transmits the frame. If the transmission fails and the ACK is not received from the receiver, then the node moves to the next backoff stage ($i+1$) and repeats the procedure. Fig. 4.1 describes the backoff procedure for two nodes, A and B, that can hear each other.

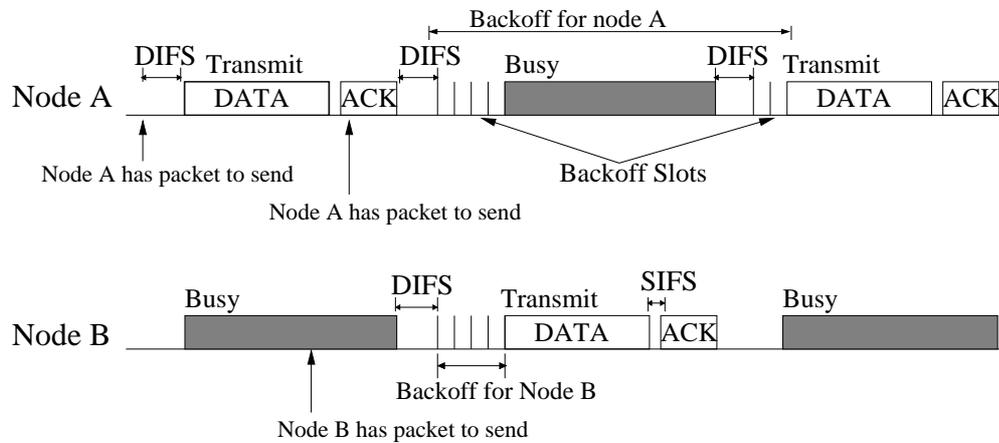


Figure 4.1: CSMA/CA: Backoff mechanism

4.3.2 Introducing History into CSMA/CA

CSMA/CA with history incorporates the history of the recent transmission activity in the medium into the conventional CSMA/CA mechanism for packet transmissions. The primary difference from conventional CSMA/CA is that the medium activity history information is used to infer the time windows that are occupied by periodic real-time flows. Time windows with periodic transmissions are considered *reserved* by the existing real-time flows, and hence there is no contention for these time windows. We assume that all the real-time flows are periodic with a common period. The most relevant example of such flows is voice applications, which generate periodic data based on the codec used (e.g., ITU-T G.711 codecs periodically generate fixed sized packets at a data rate of 64Kbps).

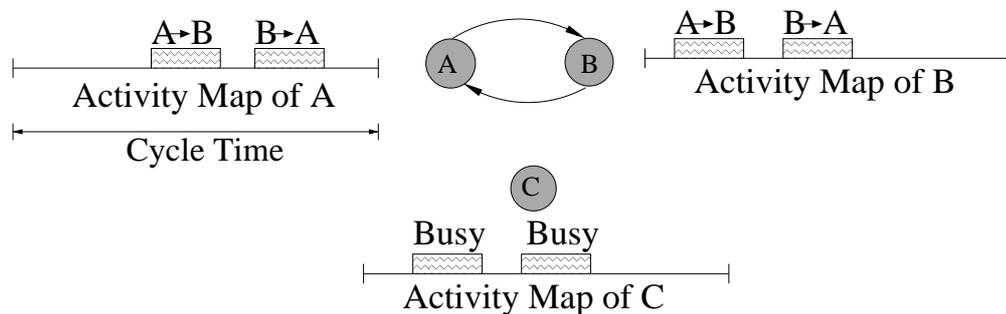


Figure 4.2: Activity Map: Nodes A and B have a bidirectional real-time flow between them. Node C is a neighbor of nodes A and B. The activity maps of nodes A, B and C contain reservations for the real-time flow between nodes A and B. Note that the cycles need not start at the same time.

We define a *cycle* as the time period between two consecutive packet transmissions of any ongoing real-time session. Therefore, the transmission activity over the last few cycles can be used to infer the real-time flow reservations of the medium. All the nodes in the network maintain a *medium activity map* containing the medium busy/free state information over a fixed number of past cycles. A node invokes conventional CSMA/CA to find a transmit opportunity, taking into account its medium activity map to ensure that it does not interfere with any existing real-time flow. A backoff mechanism similar to IEEE 802.11e, but with smaller contention window parameters, is followed to avoid collisions. As an example, consider a three node network with a bi-directional real-time flow between nodes A and B. Fig. 4.2 shows the activity maps of the medium at nodes A, B and C, assuming that node C can hear nodes A and B.

4.3.3 Sticky CSMA/CA

We view continuous time in terms of time-slots. As mentioned before, we assume that the cycle time is the same for all the nodes in the network. Note that since we assume an asynchronous framework, the cycles can start at different times at different nodes. A node with a new voice session practices a greedy algorithm to grab a free time window. A free time window is defined as a set of contiguous time-slots during which the medium is free. If the attempt to establish a flow is successful (we describe the mechanism of establishing a call later in this section), the node sticks to the time window. This means that the node sends packets periodically during the same time window in every cycle, unless the session performance, measured in terms of packet loss, degrades beyond a certain threshold. All new voice connection setup attempts and other packet transmissions honor the time windows occupied by the existing voice connections and do not contend for those slots. We describe the carrier sensing and contention mechanisms of Sticky CSMA/CA, and discuss various implementation considerations next.

Carrier Sensing Mechanism

All nodes in the network maintain a set of *Carrier Sense History Tables (CSH Tables)*. A single CSH Table is a log of medium transmissions sensed by a node

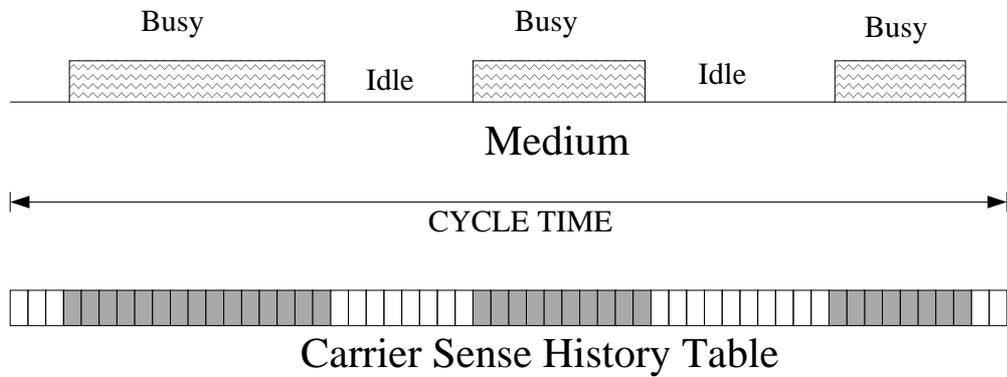


Figure 4.3: Nodes use CSMA to monitor the medium and mark their current Carrier Sense History Table according to the sensed medium busy and idle times.

over a cycle. The CSH Table contains one entry for each time-slot of the cycle. Each entry in the CSH Table is marked busy or free, based on the transmissions sensed in the medium during the corresponding time-slot. A node marks a slot as busy if it senses a transmission in its neighborhood during that time-slot, or if it transmits or receives in that time-slot. The CSH Table is updated at every medium transition among the four possible physical layer states: Idle, Sensing, Receiving, and Transmitting. It is important to note that the MAC layer does not need to be explicitly informed about the state of each time slot. MAC layer can deduce the slot states from the medium state transition information passed from the PHY layer without any additional overhead. Medium state transition information is required even in conventional CSMA MAC schemes. Fig. 4.3 shows

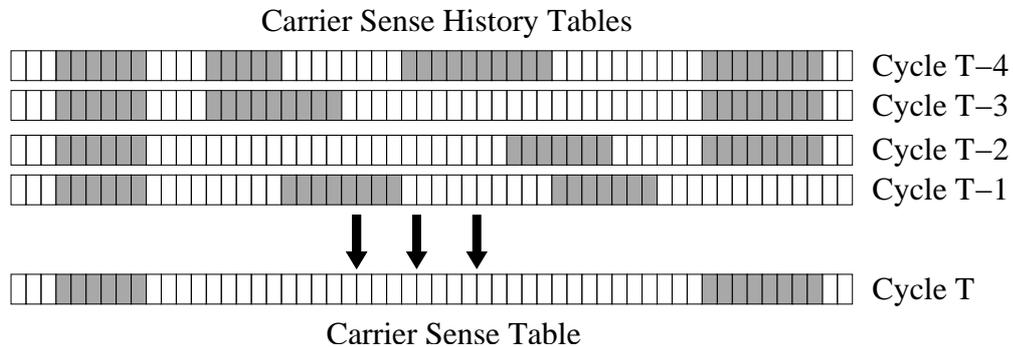


Figure 4.4: At the beginning of each cycle, nodes derive their Carrier Sense Table as a summary of their Carrier Sense History Tables. The Carrier Sense Table helps to identify periodic real-time flows.

an example of CSH Table maintenance. The number of CSH Tables required in order to maintain the history of medium activity over the past N cycles is N .

In addition to the CSH Tables, the nodes maintain another table called the *Carrier Sense Table (CS Table)*. The CS Table is an activity map of the medium that represents the predicted busy and free time-slots over the current cycle. The CS Table has a format similar to a single CSH Table. The CS Table is derived at the beginning of every cycle using the CSH Tables over the previous N cycles. A simple approach to obtain a CS Table is to do a majority decision over each slot in a cycle, e.g., if a slot is found busy in 75% of the CSH Tables over the last N cycles, it is marked as busy in the CS Table. Fig. 4.4 shows an example of CSH and CS Tables.

The CS Table is maintained only on the basis of periodicity of the sensed transmissions, and the information about the flows for which the node itself is the sender or receiver. This approach significantly reduces the control overhead as compared to the other techniques involving message exchanges, piggybacking control information over data packets, or high power transmissions. There is no extra message exchange required in this technique.

At high medium utilization, a few large sized data packets (e.g., FTP packets of size 1500 bytes) transmitted in the carrier sensing range (CSR) of a node over contiguous cycles might cause the node to falsely interpret these transmissions as voice connections. This can cause the node to mark the corresponding slots in its CS Table as busy, resulting in underutilization of bandwidth. We observe that a high percentage of data packets in the network are significantly larger than voice packets. We use this property to distinguish such data packets from voice packets. Whenever a node senses a transmission for a time interval significantly longer than that possible for a voice packet (e.g., more than double the time required for a voice packet), it concludes that it has sensed a low-priority data packet transmission and does not mark its CS Table. Also, a node does not mark its CS Table whenever it receives or overhears low-priority data packet transmissions. In addition, simple sanity checks of the CS Table data can also help in removing irrelevant history information from the CS Table. The underlying reason for not accounting for low

priority data transmissions in the CSH Table is to ensure that a node considers only relevant high priority periodic data transmission information while preparing its CS Table, thus foregoing transmission opportunities only for the periodic high priority traffic in the network. These mechanisms improve medium utilization and aid in prioritizing voice flows in the entire CSR of a node.

Contention Mechanism for Real-time Flows

To set up a new voice connection, a node checks its CS Table to find a set of contiguous free time-slots to accommodate the new flow. If a free time window is found, the node waits until the beginning of the free time window. It then employs the IEEE 802.11 CSMA/CA mechanism to transmit a Real-time RTS (R-RTS) message to the intended receiver. The R-RTS message contains the required packet transmission duration information for the new flow and a flow-identifier. The flow-identifier, along with the source node address, uniquely identifies a flow in the network. On receiving an R-RTS message, the receiver checks its CS Table (activity map) to determine whether it can support the flow from the time instant when the R-RTS transmission started (the propagation delay is negligible as compared to the transmission duration, and thus is not considered). If the corresponding time-slots are free, the node sends a Real-time CTS (R-CTS) message to the sender. The R-CTS message contains the flow-identifier and the duration

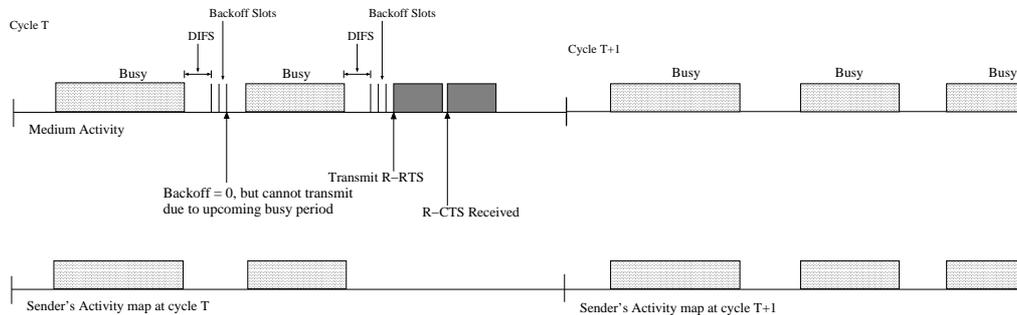


Figure 4.5: Establishing a new flow: The sender uses its activity map to select a suitable time window when it can establish a new real-time flow. The node transmits an R-RTS message at the chosen slot and receives an R-CTS message in response, thus setting up the flow.

information obtained from the R-RTS message. The sender and the receiver mark their CS Table and overwrite all the CSH Tables to indicate the presence of the new flow. The neighboring nodes of the sender and the receiver overhear the R-RTS and R-CTS messages and use the duration information to update their CSH Tables in a similar manner. The nodes in the carrier sense range (CSR) of the sender and the receiver sense this periodic transmission and update their CS Tables after a few cycles. Fig. 4.5 shows the flow setup procedure for a node.

To ensure that real-time flows win over low priority datagrams while contending for the medium during flow setup, we use an approach similar to IEEE 802.11e. IEEE 802.11e uses the contention window parameters (CW_{min} , CW_{max}) and the AIFS to prioritize different Traffic Classes (TC). Smaller CW_{min} and CW_{max} values and a shorter AIFS period ensure that a packet that belongs to a high priority

TC accesses the medium ahead of a packet belonging to a low priority TC, thereby prioritizing different TCs.

After a successful flow setup, the sender locks on to the slots; i.e., it transmits voice packets during the same time-slots in every cycle. The receiver does not send ACKs for the voice packets. ACKs result in a significant overhead for real-time flows with short packets such as voice because the size of the ACK packets is comparable to the size of the data packets [91]. Also, the physical layer preamble and header (PLCP preamble and header), which are transmitted at the lowest data rate, consume most of the transmission time required for these short packets.

Packet loss information summaries are sent by the receiving node to the sender through a *feedback mechanism*. The feedback mechanism is required for the sender to evaluate the condition of a real-time flow in terms of packet loss in the medium. This is necessary because the ACK mechanism of IEEE 802.11 is not used for real-time packets. Feedback messages are used to inform the sender of the cumulative packet losses seen by the receiver over the last M cycles. A sender node requests a feedback message by setting the Feedback Requested bit in the MAC frame header of the voice packet. The sender requests a feedback message after every $M/2$ cycles without feedback and continues sending the request with every packet of the flow until it receives a feedback message. If the maximum number of cycles M without the feedback message is reached, the sender pauses the flow and reattempts flow

setup at a different time-slot through an R-RTS/R-CTS dialog. We call this mechanism *slot reconfiguration*. The sender also utilizes the cumulative packet loss information sent in the feedback messages to decide whether slot reconfiguration is required.

Sticky CSMA/CA addresses the hidden terminal problem in the following ways. The control message exchange for slot reservation (R-RTS and R-CTS) between the sender and the receiver incorporates the receiver's view of the network before choosing slots for data transfer. A receiver does not respond to the sender's R-RTS message if it is aware of periodic transmissions over a portion of the time window requested, or if it senses the medium to be busy during these slots. In case a sender-receiver pair observes persistent packet losses over the reserved slots because of interference (e.g., from a hidden terminal), the node pair attempts to shift to another free time window that has a higher probability of being interference free. Thus, nodes avoid becoming trapped in lossy transmission patterns by attempting slot reconfiguration after the observed packet loss exceeds a certain threshold. Our performance evaluation shows that slot reconfiguration is an effective technique to avoid lossy transmission patterns.

Another effective solution to counter packet losses due to interference is *out-of-stream* packet retransmission. When the receiver of a periodic flow does not receive an expected packet, it sends a *retransmit request* to the sender. On re-

ceiving a retransmit request, the sender retransmits the corresponding packet as a *high priority* datagram. The requested packet should be sent out-of-stream because inclusion of this packet in the periodic stream will delay the subsequent packets in the stream. Our performance evaluation indicates that out-of-stream retransmissions effectively counter losses of high priority voice packets. Another approach for minimizing packet losses due to interference is to provide redundancy in the transmitted payload by including multiple voice payloads in one frame, so that the loss of a few frames does not affect the performance [91].

In general, the hidden terminal problem can be partially addressed by using sensitive antennas (i.e., with low carrier sense threshold) at the nodes so that any node that wishes to transmit can sense whether it will interfere with its neighbors' ongoing packet receptions, and thereby avoid such interference. This solution has an associated tradeoff of reduced spatial reuse because of the aggravated exposed terminal problem. Nevertheless, this approach has been proposed in literature for reducing packet collisions in systems where the required antenna sensitivity is not a design constraint [92]. We have observed this tradeoff for Sticky CSMA/CA in our performance evaluation.

Contention Mechanism for Delay-insensitive Flows

Delay-insensitive flows are assigned a lower priority than real-time flows. MAC layer operation for delay-insensitive data flows is based on CSMA/CA with history considerations. The data packet transmissions honor the reservations made by real-time flows. Nodes do not attempt to transmit data packets over slots in which high priority periodic transmission is expected. As mentioned before, data packets have higher backoff parameters and a higher AIFS than that of real-time flows to ensure that real-time flow setup requests receive a higher priority. However, the CW_{min} and CW_{max} parameters are set to lower values as compared to IEEE 802.11. This is because the time gaps left by the real-time flows might not be sufficiently large to incorporate the probable long delays due to the exponential backoff mechanism used in IEEE 802.11. Moreover, with reduced contention in a voice/data network, large contention windows are not required for collision avoidance. Our performance evaluation shows that even with smaller contention window parameters, the throughput achieved by Sticky CSMA/CA for delay-insensitive best effort traffic is comparable to that for IEEE 802.11.

With a large number of voice calls in the network, the time gaps left might not be sufficient to transmit data packets, which can lead to starvation of data flows. In systems where starvation of data flows is not acceptable, this problem can be alleviated by setting a threshold on the number of voice calls admitted. Admission

control schemes to enforce such thresholds are higher layer functionalities that are beyond the scope of this chapter. We therefore do not impose any limits on the number of voice calls in our performance evaluation.

Implementation Considerations

We discuss some important considerations for the implementation of Sticky CSMA/CA in this section.

Time-slots: Time-slots are the time units in terms of which we define the entries (marked busy or free) in the CS Table and the CSH Tables. We define time-slots of a fine granularity. We mark a time-slot as busy even when the medium is busy for only a fraction of the time-slot duration. Thus, a time-slot should be small enough so that the portion of time for which the medium is not utilized is insignificant. Also, a small slot width increases the memory requirements for the tables. Based on these considerations, we chose the time-slot duration for Sticky CSMA/CA as $20\mu\text{s}$, which is the default slot duration (used for backoff) of IEEE 802.11b. In our simulations, we assume the cycle time as 20ms (the packetization interval for voice). To allow for deviations of clocks of different nodes, extra slots (called *Leeway slots*) at the beginning and the end of the transmission are marked as busy in the CSH Table.

CS Table Maintenance: The CS and the CSH Tables are implemented as bitmaps, with each bit corresponding to a time-slot. The CSH Table for the current cycle is updated only on medium transitions in order to avoid the overhead of updating the table at every time-slot. Given a fixed cycle time, the size of the CS Table and the CSH Tables depends on the width of the time-slots. Also, the total size of the CSH Tables depends on the number of previous cycles (N) needed to derive the CS Table. For example, with the above mentioned slot width and cycle time, the memory requirement for a system with CSH Tables spanning six cycles is 875 bytes.

Retry Timer: It is necessary that the CS Table represent an accurate prediction of the transmissions in the current cycle. A node that senses the medium as free might not be allowed to contend for the channel because its CS Table indicates an impending transmission. An example scenario where this can happen is when a real-time flow ends in the carrier sensing range (CSR) of a node; the node will have the corresponding slots in its CS Table marked as busy in the current (and possibly the next) cycle, depending on the majority decision rule used to decide the busy/free status of each slot. In order to improve medium utilization in such cases, we use a *Retry timer*. A node starts a retry timer when it is not allowed to set up a new flow or transmit a data packet because the medium is expected to be busy as per its CS Table. The timer is set to expire at one of

the intermediate slots of the expected busy period. When the reattempt timer expires, the node determines whether the medium is actually busy. If the medium is found to be free, the node initiates the Sticky CSMA/CA packet transmission procedure again. This mechanism increases medium utilization even in the less probable case of an inconsistent CS Table.

Dummy Reservation Packets: When there is a packet loss at an upstream node during an ongoing multihop voice flow, the downstream nodes of the flow do not have packets to send to their next hop neighbors. In this situation, a downstream node sends dummy packets at its scheduled transmission slot to the next hop node. This transmission helps in avoiding contention for the slot by the CSR neighbors, in case of loss of a few packets (less than the loss threshold M , at which we delete the flow). Dummy reservation packets can also be used to maintain slot reservations during short silence periods (up to a maximum time limit) when VoIP packets are not generated because of the silence suppression methods being used in the codec.

Probabilistic CSH Table Updates: A node that receives a data packet from its neighbor cannot simultaneously sense the real-time transmissions of another non-interfering node in its CSR. The repeated occurrence of this event (due to high data traffic) can cause the nodes to lose the periodic flow information in their CS Tables. This in turn allows them to transmit in the slots reserved by their

CSR neighbors and cause interference. For this reason, we increase the probability of a node's retention of the real-time flow information in its CS Table by using *probabilistic updates* to the CSH Table. Whenever a node loses the periodic transmission information of its CSR neighbors from its CS Table due to receiving a long data packet, it copies the information from its CS Table to the current CSH Table with a high probability, assuming that the CS Table has more reliable information for those slots. Probabilistic CSH Table updates reduce the chances of losing the slot reservation information and result in improved performance for voice flows in the CSR of a node. We incorporate probabilistic updates in our simulations with a high probability value ($p = 0.95$).

4.4 Performance Analysis for a Clique

We now provide an approximate analytical comparison of the number of VoIP calls supported by Sticky CSMA/CA against those supported by IEEE 802.11b DCF [93] and IEEE 802.11e EDCF [62]. For this purpose, we consider single hop flows over a clique, which is a network where all nodes can hear each other. We assume that there are no hidden terminals and the radio link is error free. Despite these simplifications, the clique-based analysis can be used to predict performance for multihop flows over more complex networks, which can be roughly

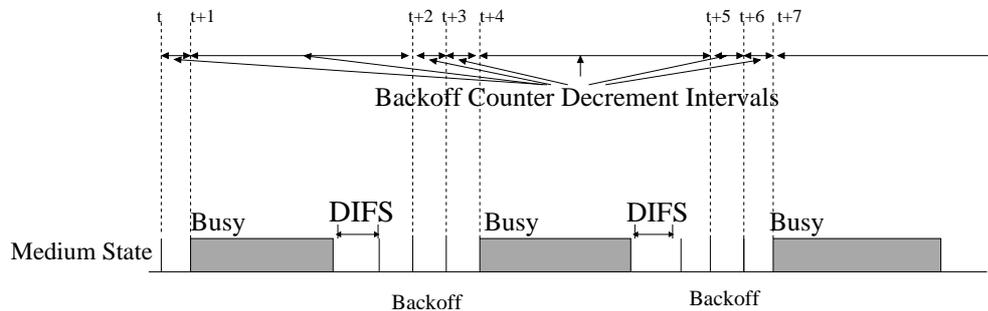


Figure 4.6: Variable time-intervals between two consecutive backoff counter decrements at an IEEE 802.11 node.

viewed as multiple cliques *stitched* together. Because of the shared nature of the wireless medium, the bandwidth required by a multihop flow can be estimated as the number of hops traveled times the bandwidth required over one link. Such computations can then be used to identify bottleneck links, which limit the maximum throughput in a multihop wireless network [94]. Applying the clique-based estimates to the nodes surrounding the bottleneck link then provides capacity estimates that match closely with the simulated performance.

4.4.1 IEEE 802.11b DCF Performance

The performance of IEEE 802.11 DCF has been analyzed in detail in [95] and [96] for single hop scenarios. References [97] and [98] incorporate the considerations for multihop flows. Our objective is to obtain the total voice call carrying capacity for an IEEE 802.11b network. We estimate the voice carrying capacity

of IEEE 802.11b using the saturation throughput analysis in [95], where saturation throughput is defined as the maximum stable throughput of the network. This analysis assumes ideal channel conditions (no hidden terminals, error free radio link, and backlogged packet queues at every node) for a network of n nodes. Through equations (4.1) to (4.3), we briefly outline the derivation found in [95] for convenience (refer to [95] for details).

Continuous time is divided into variable length time intervals, defined as *epochs*, at the beginning of which the backoff counter of each node decreases by one. These epochs can be the IEEE 802.11b *slot width* (hence forth called δ) of $20\mu s$ if the medium is free, or can also include a medium busy time interval, since the backoff counter is frozen when the medium is sensed busy. In the following analysis of IEEE 802.11b DCF, by the term epoch, we will refer to this definition of variable length time intervals. Fig. 4.6 shows the epochs as the time intervals between two consecutive backoff counter decrements. A key assumption in this model is that at each transmission attempt, irrespective of the backoff state of the node, the packet transmitted collides with a constant and independent probability p , called the *conditional collision probability*. Under these assumptions, the stochastic processes representing the backoff stage $i \in (0, m)$; and the backoff counter $k \in (0, W_i - 1)$, where $W_i = 2^i W_{min}$; can together be modeled as a bidimensional discrete-time Markov chain. Using this Markov chain model, the

Overhead	Bytes	Time (μs)
RTP	12	$12 \cdot 8 / R$
UDP	8	$8 \cdot 8 / R$
IP	20	$20 \cdot 8 / R$
IEEE 802.11 MAC	28	$28 \cdot 8 / R$
PHY Preamble		$72 / 1$
PHY Header		$48 / 2$
SIFS		20
DIFS		50
ACK	14	$14 \cdot 8 / R + \text{PHY Overhead}$
Feedback	20	$20 \cdot 8 / R + \text{PHY Overhead}$
Propagation (Δ)		1

Table 4.1: Overhead in terms of time, with average data rate R .

probability that a node transmits in a randomly chosen epoch is obtained as

$$p_t = \frac{2(1 - 2p)}{(1 - 2p)(W_{min} + 1) + pW_{min}(1 - (2p)^m)}, \quad (4.1)$$

where W_{min} is the minimum contention window, and m is the maximum allowed backoff state (i.e., $W_{max} = 2^m W_{min}$). In order to evaluate the conditional collision probability p , due to the independence assumption, p can be defined as

$$p = 1 - (1 - p_t)^{n-1}. \quad (4.2)$$

This means that the probability of collision for a transmitted frame is the probability that at least one of the other $n - 1$ nodes also transmits in the same epoch. Equations (4.1) and (4.2) can be solved for different values of n , m , and W_{min} to obtain p and p_t for these cases.

Define the probability that there is a transmission in the considered epoch as P_{tx} , and the probability that a transmission that occurs in an epoch is successful as

P_s . The normalized saturation throughput S is defined as the fraction of time the channel is used to successfully transmit payload bits, and can be evaluated as the ratio of the average time spent in the transmission of useful payload information over an average epoch duration. Define $T_{payload}$ as the time taken to transmit a payload of a fixed size, T_{tx} as the time taken for successful transmission of a complete packet including headers and other overheads, and $T_{collision}$ as the epoch duration when the medium is sensed busy by each node during a collision.

Therefore,

$$S = \frac{P_{tx}P_sT_{payload}}{(1 - P_{tx})\delta + P_{tx}P_sT_{tx} + P_{tx}(1 - P_s)T_{collision}}. \quad (4.3)$$

The normalized saturation throughput S can be evaluated for different sets of n , m , and W_{min} . The capacity of a bottleneck link in a multihop network can be approximated through S evaluated for the two node case using (4.3), where both nodes always have packets to transmit. The bottleneck link capacity constrains the maximum number of multihop flows in the network where all the nodes can hear each other. For example, if the bottleneck link supports a maximum of k voice calls, the network can support a maximum of $k/2$ two hop voice calls, or $k/3$ three hop voice calls, etc. Note that in reality, this approximation might be imprecise because all the transmitters might not transmit with the same power; the nodes might not be spaced equally apart; and the channel conditions are not

the same at all the nodes. Nevertheless, it helps to obtain an estimate of the number of calls that can be supported.

4.4.2 IEEE 802.11b VoIP Call Carrying Capacity

Table 4.1 shows the overhead associated with each voice frame sent, due to different protocol layer headers and the MAC protocol overhead. We assume that an ITU-T G.711 codec packetizing audio data every 20ms is used, which results in payloads of size 160 bytes (data rate = 64Kbps). We assume that IEEE 802.11b operates at the highest data rate of 11Mbps using the basic access mechanism (no RTS-CTS). We also assume that IEEE 802.11b uses the *short* physical layer preamble (PLCP preamble), which reduces the synchronization time from $192\mu\text{s}$ (when the long physical layer preamble is used) to $96\mu\text{s}$. The values mentioned above correspond to the state-of-the-art IEEE 802.11b VoIP equipment commercially available [83]. Define T_{frame} as the time required to transmit a frame

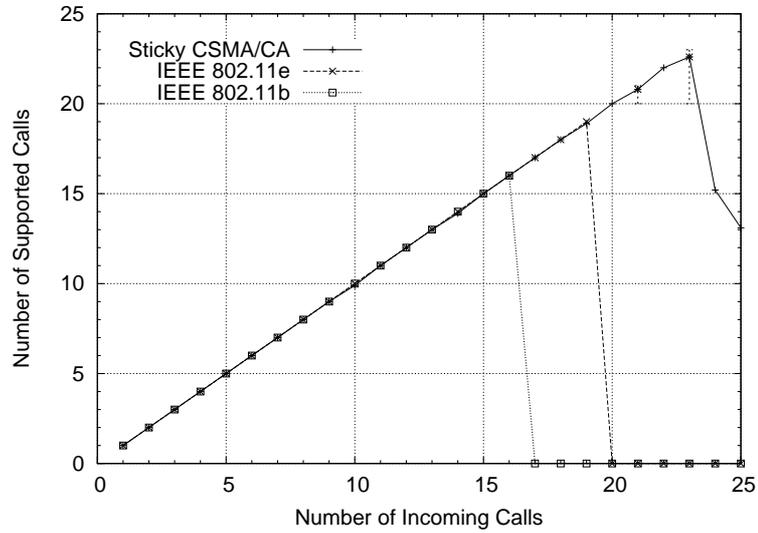


Figure 4.7: The total number of voice calls supported as a function of the incoming calls for a two node network, with nodes 150m apart.

(payload and headers), and Δ as the propagation time. From Table 4.1, we have

$$T_{payload} = \frac{160 \cdot 8}{11} \mu s,$$

$$\begin{aligned} T_{frame} &= T_{PHY} + T_{MAC} + T_{RTP+UDP+IP} + T_{payload} \\ &= 262 \mu s, \end{aligned}$$

$$\begin{aligned} T_{tx} &= T_{frame} + SIFS + ACK + DIFS + 2 \cdot \Delta \\ &= 430 \mu s, \end{aligned}$$

$$T_{collision} = T_{frame} + DIFS + \Delta = 313 \mu s.$$

Solving equations (4.1) and (4.2) numerically for the two node case ($n = 2$, $W_{min} = 32$, and $m = 5$, (as (CW_{min}, CW_{max}) for IEEE 802.11b are $(31, 1023)$) [93] and plugging in the required probability and transmission time values in (4.3), we get the normalized saturation throughput $S = 0.1924$. This gives the saturation throughput = $S * 11\text{Mbps} = 2.1164\text{Mbps}$. Each voice call is equivalent to two unidirectional constant bit rate (CBR) flows of data rate 64Kbps. Therefore, the required data rate for a single VoIP flow = 128Kbps. Thus, we obtain the maximum number of calls supported as 16. This number matches with the capacity we obtain from simulations, as shown in Fig. 4.7. The simulation parameters are summarized in Section 4.5. Note that if the long preamble is used, the maximum number of calls supported by IEEE 802.11b network decreases to 12.

The IEEE 802.11e EDCF network for a voice-only scenario is similar to IEEE 802.11b DCF, with different contention window parameters (i.e., $(7, 31)$ instead of $(31, 1023)$ for voice packets). Note that prioritization of different access categories does not play a role in this scenario because we are only considering voice calls. The total number of voice calls obtained in this case is 19, which matches with the capacity obtained via simulations (see Fig. 4.7). It is important to note that optional IEEE 802.11e protocol features (such as transmission opportunity (TXOP) that facilitates contention free bursting) when enabled, result in a differ-

ent overall network behavior. The effect of TXOP has been accounted for in the simulation based performance evaluation in Section 4.5.

4.4.3 Sticky CSMA/CA VoIP Call Carrying Capacity

Sticky CSMA/CA requires the nodes with real-time flows to stick to a periodic transmission schedule. The other nodes in the carrier sensing range (CSR) neighborhood that hear the periodic transmissions do not attempt to contend for the medium during the slots in which they expect a periodic transmission. By sticking to a periodic transmission schedule and using contention only to *reserve* slots for a real-time flow, Sticky CSMA/CA reduces the number of contending nodes in the medium drastically (in contrast with IEEE 802.11, where the number of contenders at any instant equals the number of nodes that have packets to transmit). This leads to substantial gains in terms of reducing the overhead involved with backoff for every packet transmitted in the medium. Based on this fact, we reason that in a network with only voice flows, the total number of contenders at any time instant in a Sticky CSMA/CA network equals the number of new arriving calls (and the calls being reconfigured to a different time window), which is significantly less than the total number of nodes with packets to transmit at any instant. With this observation, we argue that it is reasonable to assume that, on average, the first stage of backoff yields a winner with a high probability, even with

a very small CW_{min} . Thus, we use contention window parameters smaller than IEEE 802.11b for real-time flow setup requests in a Sticky CSMA/CA network: $CW_{min} = 3$ and $CW_{max} = 7$. These contention window parameters are used in a manner similar to IEEE 802.11e for service differentiation. We do not need to account for the slots wasted due to collisions in the calculation of the maximum number of VoIP connections because the time-slots wasted in the current cycle due to a collision (and the consequent DIFS interval and backoff) can effectively be used by a new VoIP call in the coming cycles.

Sticky CSMA/CA uses feedback messages sent periodically or on-demand to the sender in place of ACK messages for voice packets. In our simulations, the sender demands a feedback message every sixth cycle. Thus, the average additional medium access time required per cycle to transmit feedback messages for a real-time flow is one-sixth of the time required for transmitting a single feedback message. Note that the slots reserved for a voice packet also include leeway slots ($L = 1$) on both sides of the reserved slots.

We express all the time requirements in terms of time-slots (with slot size $\delta = 20\mu s$). Because a cycle spans 20ms (the rate at which the 160 byte voice packets are generated), $T_{cycle}^s = 1000$. The minimum average gap between the slots reserved for a new voice call placed next to an already existing reservation in the cycle is $DIFS + E[x] \cdot \delta$, where $E[x]$ is defined as the expected value of a

random variable x that models the portion of the time gap (expressed in terms of δ) between an established voice call and a new call placed next to it that is contributed by the backoff duration associated with every new call. The last reserved slot of a new call need not have any time gap from the first slot of an already existing call reservation placed next to it. This means that two adjacent calls might or might not have any time gap between them, depending on the order of their arrival. We assign equal probabilities to the two possible call arrival orders. Thus, we define the average minimum time gap associated with each call as $T_{empty}^s = 0.5(\frac{DIFS}{\delta} + E[x])$, where $E[x] = \frac{CW_{min}}{2} = 1.5$ slots. Define the total slots required for transmitting a frame, including the overheads as T_{tx}^s ; total slots required for feedback messages (of size 20 bytes) as $T_{feedback}^s$; and the total number of slots required as $Slots_{reqd}$. Then,

$$\begin{aligned}
 T_{empty}^s &= \frac{(E[x] \cdot \delta + DIFS)}{2\delta} = 2, \\
 T_{tx}^s &= \frac{(T_{frame} + \Delta)}{\delta} + T_{empty}^s = 16, \\
 T_{feedback}^s &= \frac{T_{PHY} + \frac{20 \cdot 8}{11} + \Delta + DIFS}{\delta} + E[x] \\
 &= 10, \\
 Slots_{reqd} &= (T_{tx}^s + \frac{T_{feedback}^s}{6}) + 2 \cdot L = 20.
 \end{aligned}$$

Therefore, the number of slots required per voice connection comprising two unidirectional flows is $T_{voip}^s = 2 * Slots_{reqd} = 40$. Thus, the maximum number of

calls that can be supported is estimated as $T_{cycle}^s/T_{voip}^s = 25$. This analysis shows that Sticky CSMA/CA has more than 50% additional call carrying capacity as compared to IEEE 802.11b, and more than 25% additional call carrying capacity as compared to IEEE 802.11e. The maximum number of calls supported for the two node scenario obtained by simulations is 23 (see Fig. 4.7). The error bars at 21 or more incoming calls indicate that the number of supported calls varies depending on call arrival times in the different simulation runs. The loss at high number of incoming calls can be attributed to the packing inefficiency while placing calls in a cycle. We do not employ any mechanism to tightly pack the calls in the cycle where a greater number of calls could be supported.

Thus, Sticky CSMA/CA achieves almost 45% more VoIP call capacity as compared to IEEE 802.11b, and 20% more VoIP call capacity as compared to IEEE 802.11e over the two node scenario. The underlying reasons are the gains due to periodic scheduling and reduced contention as discussed earlier.

While evaluating the voice call carrying capacity, we have considered only voice traffic in the network. When data is introduced, Sticky CSMA/CA performs much better than IEEE 802.11b/e in terms of the number of supported calls because of the prioritization of voice traffic induced by grabbing a periodic transmission schedule. Moreover, the lower contention window and AIFS parameters help in achieving a higher priority for voice flow setup requests, which aids the overall

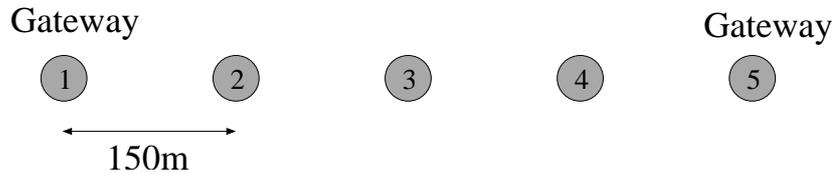


Figure 4.8: Line topology: the nodes are 150m apart. The edge nodes 1 and 5 act as gateway nodes, with all the calls being routed through them.

voice throughput. In Section 4.5 we demonstrate that even with a higher priority to voice calls and no packing optimization, the data throughput achieved by Sticky CSMA/CA is comparable to IEEE 802.11b/e.

We now evaluate the performance of Sticky CSMA/CA and compare against both IEEE 802.11b and IEEE 802.11e via simulations.

4.5 Performance Evaluation

In this section, we evaluate the performance of Sticky CSMA/CA in terms of VoIP calls supported along with the low priority data traffic, over different network scenarios. Through simulations, we demonstrate that Sticky CSMA/CA achieves a much higher VoIP call carrying capacity than IEEE 802.11b/e while providing the required QoS to the VoIP flows.

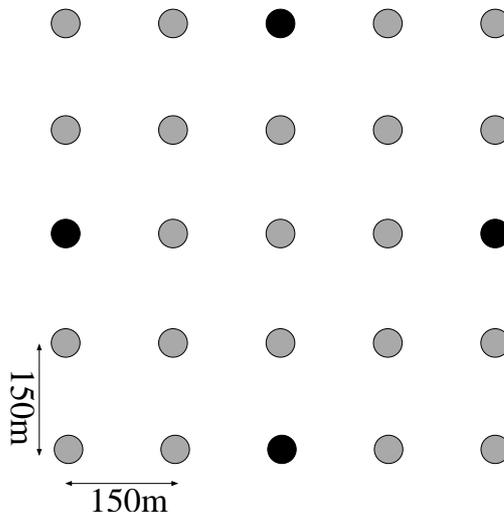


Figure 4.9: Grid topology: 25 nodes placed 150m apart in a square grid. The nodes colored black act as the gateway nodes for the network.

4.5.1 Simulation Setup

We use the QualNet Network Simulator [61] to evaluate the performance of Sticky CSMA/CA. We consider three different network topologies: a single line of five nodes, spaced 150m apart; a regular grid of 25 nodes, spaced 150m apart; and an irregular grid of 25 nodes with an average distance of 150m between neighboring nodes. Figs. 4.8 and 4.9 illustrate the line and the regular grid topologies along with the gateway nodes that connect the mesh network to the wired backbone. We discuss the irregular grid topology later in this section. These topologies model different degrees of medium contention and interference among nodes in mesh networks and aid in understanding the performance of Sticky CSMA/CA and IEEE 802.11b/e over real mesh network deployments. Because all

(a) Physical layer and MAC layer

MAC Layer Parameters	Values	PHY Layer Parameters	Values
Slot time	$20\mu s$	Propagation Model	Two Ray Model
SIFS	$10\mu s$	Antenna	Omnidirectional
DIFS	$50\mu s$	PHY Noise Factor	10
IEEE 802.11b CW_{min}	31	Transmission Power	11.0 dBm
IEEE 802.11b CW_{max}	1023	Data Rate	11.0 Mbps
		PLCP Preamble	Short Preamble (72 bits)
		Synchronization Time	$96\mu s$
		Propagation Delay	$1\mu s$
		Constant Shadowing Mean	4 dB

(b) Sticky CSMA/CA

Sticky CSMA/CA Parameters	Values
AIFS (Voice)	$50\mu s$
AIFS (Data)	$70\mu s$
(CW_{min}, CW_{max}) Voice	(3,7)
(CW_{min}, CW_{max}) High Priority Data	(15,31)
(CW_{min}, CW_{max}) Background Data	(31,63)
R-RTS	26 bytes
R-CTS	20 bytes
MAC Data Header	30 bytes
Cycle Time	20 ms
Missed Feedback Threshold	10

(c) IEEE 802.11e

IEEE 802.11e Parameters	Values
AIFS (Voice)	$50\mu s$
AIFS (Background Data)	$150\mu s$
(CW_{min}, CW_{max}) Voice	(7,15)
(CW_{min}, CW_{max}) Video	(15,31)
(CW_{min}, CW_{max}) Background Data	(31,1023)
TXOP Limit for Voice	$3008\mu s$

Table 4.2: Simulation parameters.

the (outgoing/incoming) voice calls and data traffic are routed to their destinations through the wired backbone, the gateway nodes act as terminal points for the voice and data traffic in a mesh network. This arrangement of nodes models the current single radio mesh network deployments, where the closest gateway for each mesh node is usually no more than two hops away [99]. Because Sticky CSMA/CA is designed for mesh networks, all the nodes in the network are assumed to be stationary.

We use the two-ray propagation model with constant shadowing as our physical layer model. To evaluate Sticky CSMA/CA using a simple and tractable physical layer model, we consider a flat terrain and neglect the time-varying effects of multipath fading. Table 4.1(a) lists the simulation parameters used in our simulations. The physical layer parameters used for Sticky CSMA/CA result in a transmission range (TR) of 200m, and a carrier sensing range (CSR) of approximately 450m. For nodes spaced 150m apart, the interference range (IR) is approximately 420m. We use the short physical layer preamble provided by the IEEE 802.11b standard. IEEE 802.11b uses the basic access mechanism (RTS-CTS turned off). Table 4.1(c) lists the parameters used for service differentiation in IEEE 802.11e. The simulation parameters used for IEEE 802.11b/e simulations are known default parameters that are recommended by the respective standards or are commonly used by the state-of-the-art equipment available in the market [62, 83, 93].

To simulate the bi-directional VoIP calls, we use two simultaneous constant bit rate (CBR) flows between the source and the destination nodes, one in each direction. The network layer is assumed to support priority queuing of voice packets over the background data traffic. We consider an ITU-T G.711 codec packetizing 20ms of audio data in 160 bytes of payload (data rate 64Kbps). We model call arrival as a Poisson process. Table 4.1(b) lists parameter values specific to Sticky CSMA/CA used in the simulations. In order to prioritize real-time traffic

over data packets, we set the contention window parameters (CW_{min} , CW_{max}) for voice setup requests as (3, 7), for high priority data as (15, 31), and for background data as (31, 63). The Arbitration Interframe Space (AIFS) is $50\mu s$ for voice and high priority data, and $70\mu s$ for background data traffic. These parameters ensure a higher priority to the voice traffic than the data traffic. The carrier sense history information is maintained for the past six cycles. A sender node that does not receive a feedback message from the receiver for 10 consecutive cycles (Missed Feedback Threshold) attempts slot reconfiguration. Slot reconfiguration is also attempted when the feedback message sent by the receiver indicates more losses than a certain threshold (10% to 20% in our simulations).

We evaluate the performance of Sticky CSMA/CA and compare it with IEEE 802.11b and IEEE 802.11e in terms of the total number of VoIP calls supported with the required QoS, i.e., bounded packet loss, average delay, and delay jitter. We explore the impact of background traffic (both TCP and UDP) on the call carrying capacity and estimate the achievable throughput for these flows. Note that we do not address the starvation of background data flows due to voice traffic in our simulations, since the objective is to understand the effect of these flows on the voice call traffic and vice-versa. Our results can be used in practical systems to determine the maximum number of voice calls that can be admitted, while sustaining a given data throughput. Call admission control, resource con-

strained routing, and end-to-end signaling for call establishment are higher layer functionalities that are beyond the scope of this chapter.

We impose a delay constraint of 50ms for voice packets, since this constraint allows ample time for the transfer of voice packets over the backbone network to ensure a total end-to-end delay of 150ms. We consider a VoIP call as useful only if 95% of the packets sent in both directions are delivered within the delay bound. We use these constraints to model the QoS requirements of VoIP calls [100]. The simulation results have been averaged over 10 runs.

We generate the irregular grid topology by uniformly varying the placement of nodes within a square of side 40m around the corresponding positions in the regular grid topology. This arrangement of nodes ensures that, with a high probability, neighboring nodes are within the TR of each other so that links between them still exist, but with different link qualities. We use a different uniform random grid topology with each simulation run.

4.5.2 Line Topology

Consider a mesh network that consists of a single line of five nodes as shown in Fig. 4.8. Nodes 1 and 5 act as gateways for the network. All the calls are routed out of the network through these gateways. Thus, one of the two gateways in the

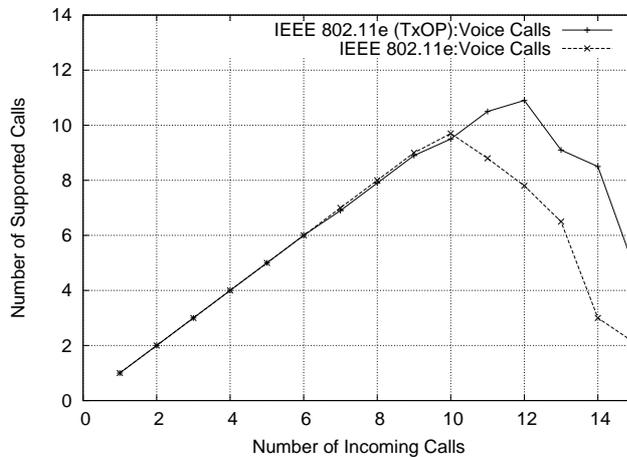
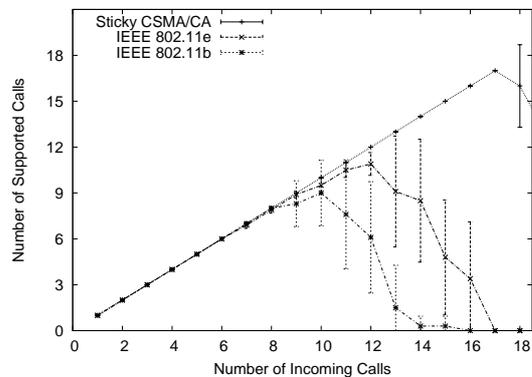


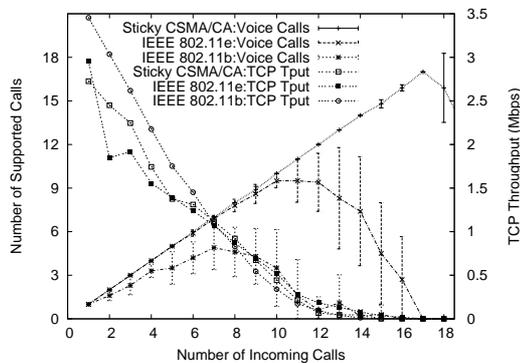
Figure 4.10: The effect of TXOP on the total number of voice calls supported over a five node line topology.

network is always the terminal point of the CBR flows corresponding to a VoIP call in the network.

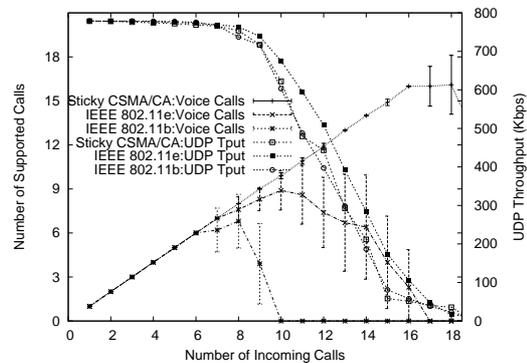
As mentioned in Section 4.2, IEEE 802.11e has an optional feature, called Transmission Opportunity (TXOP), that allows contention-free bursting for the corresponding access category (AC) after a node has won medium contention. TXOP means that a node that has won medium contention can continuously send data up to a maximum allowed transmission duration for that AC without having to contend again. As shown in Fig. 4.10, over the line topology, the total number of calls supported by IEEE 802.11e improves from 9 to 11 with TXOP. We incorporate TXOP support in our IEEE 802.11e simulations to increase the IEEE 802.11e voice call carrying capacity.



(a) Voice traffic only.



(b) Voice and TCP background traffic.



(c) Voice and UDP background traffic.

Figure 4.11: Sticky CSMA/CA call carrying capacity for the line topology. All figures show 95% confidence intervals for the average number of supported calls.

Fig. 4.11(a) shows the average number of voice calls supported as a function of the number of incoming calls. We generate voice calls by randomly picking the nodes (other than gateways) in the network and assigning the nearest gateway as the terminal node for each call. Hence, calls traverse either a single hop or multiple hops to the gateway depending on the source node. Each data point in Fig.

4.11(a) corresponds to the average number of calls supported for a fixed number of incoming calls. 95% confidence intervals for the average number of supported calls are also shown. Observe from Fig. 4.8 that Node 3 can carrier sense all the nodes in the network. Since the links from Node 3 are effectively shared by all the other nodes in the network, Node 3 is a bottleneck node. From Fig. 4.11(a), we observe that the total number of calls supported by Sticky CSMA/CA is 17, whereas IEEE 802.11e and IEEE 802.11b support a maximum of 11 and 9 calls respectively. The gains in the VoIP call carrying capacity for Sticky CSMA/CA result from reduced contention overhead, reduced average backoff time, and the elimination of per-packet ACKs. Notice that in the absence of a call admission control mechanism, the performance degrades when the number of incoming calls exceeds the total capacity of the network. Lack of resource reservation and medium contention for every packet transmission result in reduced total call carrying capacity for IEEE 802.11b/e.

Fig. 4.11(b) shows the average number of supported voice calls when there is background TCP traffic consisting of payloads of varying sizes (maximum segment size of 1460 bytes). In this evaluation scenario, we have two TCP flows in the network, originating from gateway nodes 1 and 5, and terminating at nodes 2 and 3, respectively. This TCP traffic scenario models the long-lived FTP and other TCP based content download applications from the gateway nodes that connect to the

Internet. The data traffic workload ensures that all the nodes in the network have backlogged low priority data queues, and that the nodes are directly affected by the background TCP traffic dynamics. The capacity of Sticky CSMA/CA, in terms of the total number of supported calls, remains the same even in the presence of background TCP traffic of varying payload sizes. This is because of the effective priority given to voice packets over data transmissions in the entire carrier sensing range (CSR) of a node. IEEE 802.11e (with TXOP) supports a higher number of voice calls as compared to IEEE 802.11b because of service differentiation, but the total number of calls supported by IEEE 802.11e is still significantly less than that supported by Sticky CSMA/CA. Due to interference and lack of prioritization in IEEE 802.11b, the number of supported voice calls reduces drastically with TCP traffic, since different nodes with voice or data packets contend with each other for medium access. For data traffic in Sticky CSMA/CA, the bandwidth left over by the real-time flows is sufficient to maintain average data throughput slightly lower than IEEE 802.11b but relatively higher than IEEE 802.11e.

With fewer voice calls in the network, the effect of interference caused by heavy data traffic in terms of loss of many calls for IEEE 802.11b is apparent in Fig. 4.11(b). IEEE 802.11e is able to recover from these losses by multiple fast retransmissions of voice packets aided by shorter contention windows for voice. The TXOP feature helps in recovering from delays due to retransmissions by

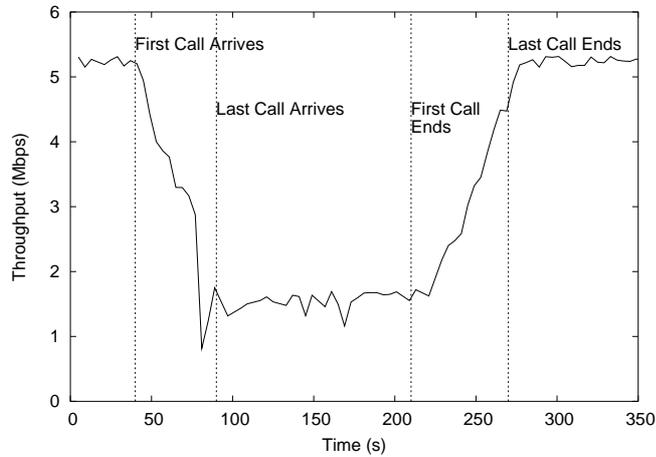


Figure 4.12: Instantaneous throughput of a FTP flow in a representative simulation for Sticky CSMA/CA.

transmitting the accumulated voice packets together in bursts. Sticky CSMA/CA effectively copes with the interference caused by data traffic with the help of techniques discussed in Section 4.3.3, which results in approximately 200% additional call carrying capacity as compared to that of IEEE 802.11b, and 70% additional call carrying capacity as compared to that of IEEE 802.11e.

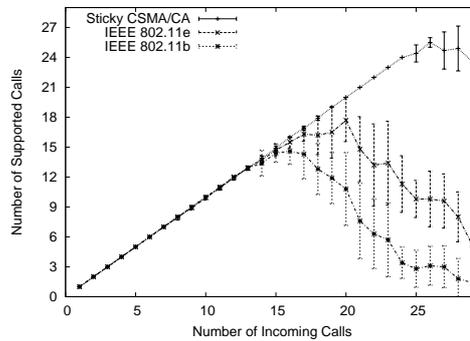
Fig. 4.11(c) shows the average number of supported voice calls in the presence of background UDP traffic. We use three UDP flows, each with a data rate of 256Kbps and a payload size of 1460 bytes. These UDP flows originate from the three non-gateway nodes in the network, and are routed to the nearest gateway node. As in the TCP background traffic scenario, this data traffic workload ensures that all the nodes in the network have low priority data packets to send, and that

they have to deal with the background UDP traffic dynamics to sustain voice calls without performance degradation. We observe from Fig. 4.11(c) that the background UDP traffic has no significant effect on Sticky CSMA/CA voice call performance. UDP throughput for Sticky CSMA/CA starts to decrease after the number of incoming calls increases to an extent that the leftover gaps between the voice calls are not sufficient to sustain the UDP flows. Sticky CSMA/CA UDP throughput is comparable to that of IEEE 802.11b/e.

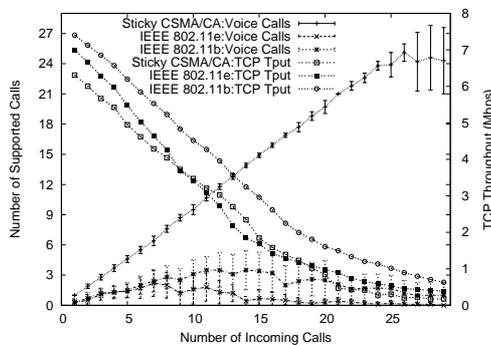
Fig. 4.12 shows the instantaneous TCP throughput for Sticky CSMA/CA with increasing number of incoming calls over the line topology. Observe that TCP throughput decreases as the number of calls increases, demonstrating that the priority mechanism provides effective service differentiation between background data traffic and voice. Data flows adapt to the voice traffic in the network depending on the leftover capacity.

4.5.3 Grid Topologies

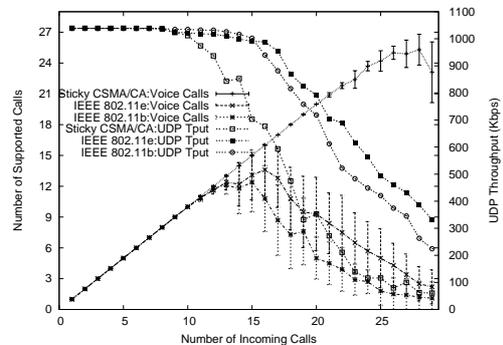
Consider a 25 node regular grid topology as shown in Fig. 4.9. All VoIP calls in the system are routed out of the network through the gateways. The gateways are also assumed to be the sources of long-lived TCP or UDP flows. Though the grid topology provides more opportunities for spatial reuse, it is also characterized by increased contention and interference among nodes. We compare the performance



(a) Voice traffic only.



(b) Voice and TCP background traffic.



(c) Voice and UDP background traffic.

Figure 4.13: Sticky CSMA/CA call carrying capacity for the regular grid topology. All figures show 95% confidence intervals for the average number of supported calls.

of Sticky CSMA/CA with IEEE 802.11b and IEEE 802.11e in terms of the total number of VoIP calls supported with the required QoS over the grid topology. We explore the impact of background traffic (both TCP and UDP) on the call carrying capacity.

Fig. 4.13(a) shows the average number of voice calls supported as a function of the number of incoming calls. We generate calls by picking nodes in the network at random and assigning the nearest gateway as the terminal node for each call. Ties between the gateways are broken arbitrarily. For UDP or TCP flows, the gateway nodes are chosen as the sources, and the destination nodes are randomly picked among the nodes serviced by the gateway. We see that the maximum number of calls supported by Sticky CSMA/CA is 25, whereas IEEE 802.11e supports a maximum of 18 calls and IEEE 802.11b supports only 15 calls. The increase in the number of calls supported by Sticky CSMA/CA and IEEE 802.11b/e over the grid topology is because of increased spatial reuse of the medium.

Fig. 4.13(b) shows the average number of supported voice calls as a function of the number of incoming calls in the presence of background TCP traffic with varying payload sizes (maximum segment size of 1460 bytes). In this scenario, we generate the background TCP traffic by initiating one TCP flow from each gateway node and choosing a random non-gateway node from its service area as the destination, routing packets via preconfigured shortest path routes. Thus, the background TCP traffic dynamics differ over different simulation runs with different seeds. We observe that the maximum call carrying capacity for Sticky CSMA/CA remains the same as that without the background TCP traffic. The effect of high interference in the grid topology is significantly more pronounced

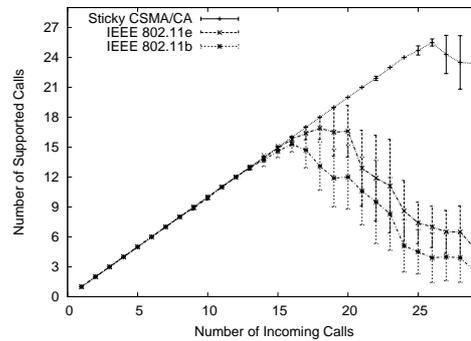
for IEEE 802.11b/e in terms of majority of incoming calls being dropped. Sticky CSMA/CA is able to cope with this increased interference except for a small probability of call drop in some cases with fewer incoming calls. As the number of VoIP flows increases, the interference problem is alleviated as the time gaps between VoIP flows decrease, reducing the overall data throughput. The voice call carrying capacity of Sticky CSMA/CA in this scenario is many times higher than that of IEEE 802.11b/e.

The increased data throughput observed for IEEE 802.11b can be attributed to the lack of prioritization of voice over background data because of which background data flows achieve much higher throughput at the expense of voice calls. This lack of service differentiation adds to the effect of increased interference in the regular grid topology and further degrades the voice call performance of IEEE 802.11b. With fewer calls in the network, TCP throughput for Sticky CSMA/CA is slightly less than that for IEEE 802.11e, but this increased throughput for IEEE 802.11e is at the expense of most of the incoming voice calls being dropped. Note that the voice call performance of Sticky CSMA/CA does not change significantly in the presence of TCP packet transmissions of varying sizes. This demonstrates that the CS Tables maintained by the nodes in the network contain the relevant periodic high priority flow information and are not corrupted by the data packet transmissions even under heavy TCP traffic.

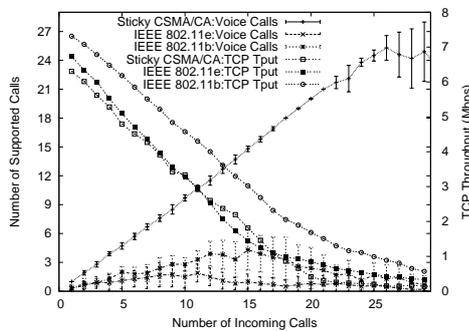
Fig. 4.13(c) shows the average number of supported calls when there is background UDP traffic in the network. The background UDP traffic is generated in a similar manner as the background TCP traffic, with each UDP flow having a data rate of 256Kbps and a payload size of 1460 bytes. Again, a capacity gain is observed. The sharp fall in data throughput as compared to IEEE 802.11b/e after 10 voice calls can be attributed to the time gaps between the voice transmissions being too small for data packets to be transmitted. Note that the data throughput is approximately equal for Sticky CSMA/CA and IEEE 802.11b/e until the number of voice calls in the network reaches the maximum that can be supported by IEEE 802.11b/e. The loss of background data for more incoming calls is the cost of supporting significantly more high priority voice calls.

Fig. 4.14 illustrates performance comparison in terms of voice call carrying capacity of the network in voice only case, and in presence of background TCP/UDP traffic over the irregular grid topology. We observe that the results obtained for the irregular grid topology and the regular grid topology show very similar trends.

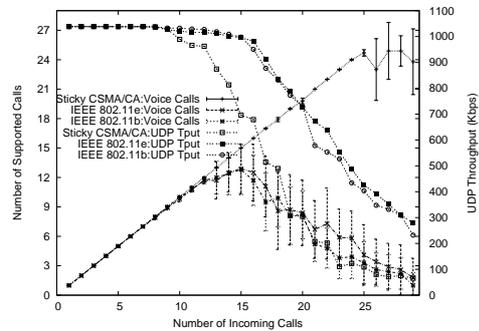
We make the following observations regarding the supported calls results for different simulation runs for all the test topologies: the variation of the results for different simulation seeds is minimal until the number of calls approaches the system's call-carrying capacity for Sticky CSMA/CA and IEEE 802.11b/e for the voice only case. A similar trend is observed for voice call results with UDP



(a) Voice traffic only.



(b) Voice and TCP background traffic.



(c) Voice and UDP background traffic.

Figure 4.14: Sticky CSMA/CA call carrying capacity for the irregular grid topology. All figures show 95% confidence intervals for the average number of supported calls.

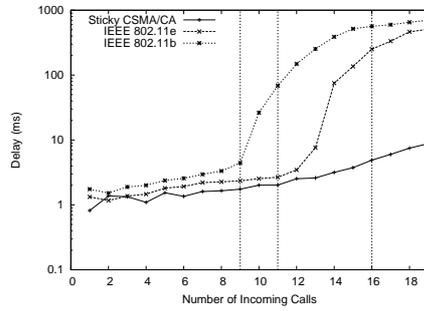
background traffic, except that the results start to vary earlier than the voice only case. For TCP background traffic, the variation of the number of supported calls is higher than that observed for the voice only and the UDP background traffic cases for all the protocols. Interestingly, for both IEEE 802.11b and 802.11e over the regular and irregular grid topologies, the most dominant result for all data points (over most seeds) is zero calls, and the deviations are non-zero numbers that

account for the very low average call carrying capacity for both IEEE 802.11b/e. This is because of high interference and contention that result in heavy packet losses for both IEEE 802.11b/e.

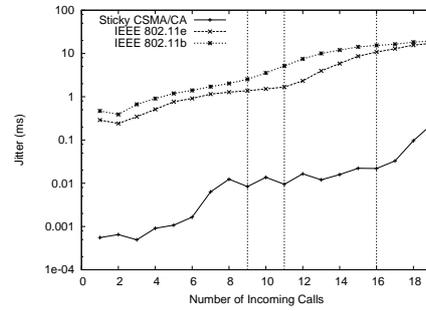
4.5.4 QoS Performance: Delay and Delay Jitter

We compare the QoS performance of Sticky CSMA/CA and IEEE 802.11b/e in terms of the average end-to-end delay and delay jitter for voice calls. Figs. 4.15, 4.16 and 4.17 illustrate the average end-to-end delay and delay jitter for Sticky CSMA/CA and IEEE 802.11b/e for the line topology, the regular grid topology and the irregular grid topologies, respectively. The traffic in the network consists of only voice calls. The delay performance for Sticky CSMA/CA remains fairly constant with increasing number of voice calls. We notice the steep rise in delay for IEEE 802.11b and IEEE 802.11e when the number of incoming calls exceeds the network capacity. Because of the TDM-like synchronization achieved by Sticky CSMA/CA, the average delay jitter is orders of magnitude less compared to that obtained for IEEE 802.11b/e.

Performance trends of Sticky CSMA/CA in terms of the voice call carrying capacity and the offered QoS are very similar over the line topology, the regular grid topology and the irregular grid topologies. This observation implies that the performance of Sticky CSMA/CA scales over more general, irregular mesh

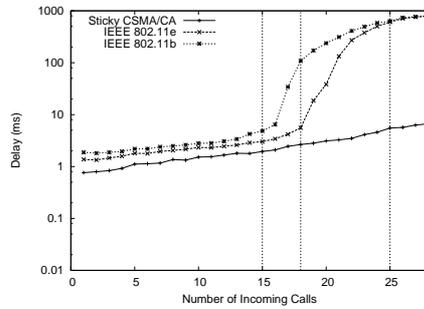


(a) End-to-End Delay.

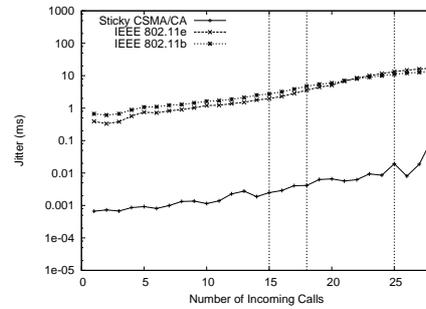


(b) Delay Jitter.

Figure 4.15: Line topology: QoS performance for voice calls.

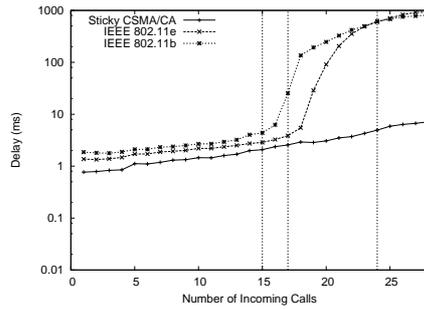


(a) End-to-End Delay.

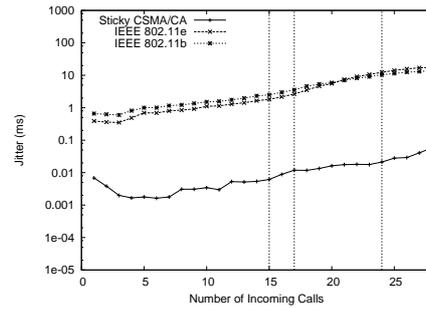


(b) Delay Jitter.

Figure 4.16: Regular grid topology: QoS performance for voice calls.



(a) End-to-End Delay.



(b) Delay Jitter

Figure 4.17: Irregular grid topology: QoS performance for voice calls.



Figure 4.18: Nine node line topology: the edge nodes 1 and 9 act as gateway nodes.

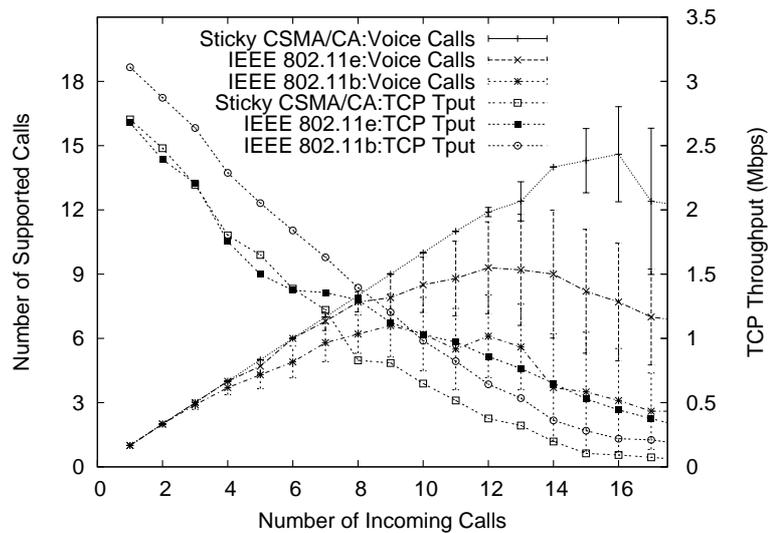


Figure 4.19: Voice call support for nine node single line topology network.

network topologies. In order to gain more insight into Sticky CSMA/CA performance over different practical mesh network scenarios, we next study the effect of increasing number of hops on Sticky CSMA/CA and IEEE 802.11b/e voice call performance. We then look at the voice call performance over networks employing silence suppression to optimize bandwidth consumption of voice calls.

4.5.5 Effect of Increasing Number of Hops

In order to evaluate the effect of paths with large number of hops on Sticky CSMA/CA voice call performance, we consider a mesh network with nine nodes arranged in a single-line topology (see Fig.4.18). We include background TCP traffic to model realistic network traffic loads. The voice call and data traffic load generation methodology is similar to that used for the previous evaluations in this section. Thus, there are voice calls that traverse two, three or four hops over the network depending on the nodes chosen as sources of voice calls for each simulation run. Note that for calls traversing four hops, the gateway node is out of the carrier sense range of the source node, and thus it does not have complete carrier sensing information of the VoIP call in its CS Table. There are two TCP flows in the network which traverse three hops (with gateway nodes 1 and 9 as sources and nodes 4 and 6 as destinations). Fig. 4.19 compares the total number of supported calls and aggregate TCP throughput for Sticky CSMA/CA and IEEE 802.11b/e. We observe that Sticky CSMA/CA has a higher call carrying capacity in this scenario. This is consistent with the results reported for other network topologies, and shows that Sticky CSMA/CA performs better than IEEE 802.11b/e even for paths with more hops.

4.5.6 Silence Suppression

Silence suppression is an important configurable feature that is used to reduce the bandwidth usage of VoIP flows. End nodes that employ silence suppression do not send packets during the silence periods in the VoIP call thereby saving bandwidth. Many common audio codec standards along with the Real-time Transport Protocol (RTP) support silence suppression [101]. We compare the performance of Sticky CSMA/CA with IEEE 802.11b/e protocols over an example scenario of the regular grid topology with background TCP traffic. We use Brady's On-Off Markov Modulated Fluid model [102] to model voice calls with talk spurts and silence periods. Brady's model has been commonly used in the VoIP literature because of its simplicity and ability to capture the dynamics of telephonic conversations with reasonable accuracy. In this model, the silence periods and the talk spurt periods of a voice source are exponentially distributed random variables with means of $1.35S$ and $1S$, respectively. The probability p that the sending end is active is 0.43.

Fig. 4.20 compares the voice call carrying capacity of Sticky CSMA/CA with IEEE 802.11b/e when silence suppression is used. We observe that Sticky CSMA/CA supports a higher number of voice calls. The underlying reason for this performance gain is that Sticky CSMA/CA is able to exploit the predictabil-

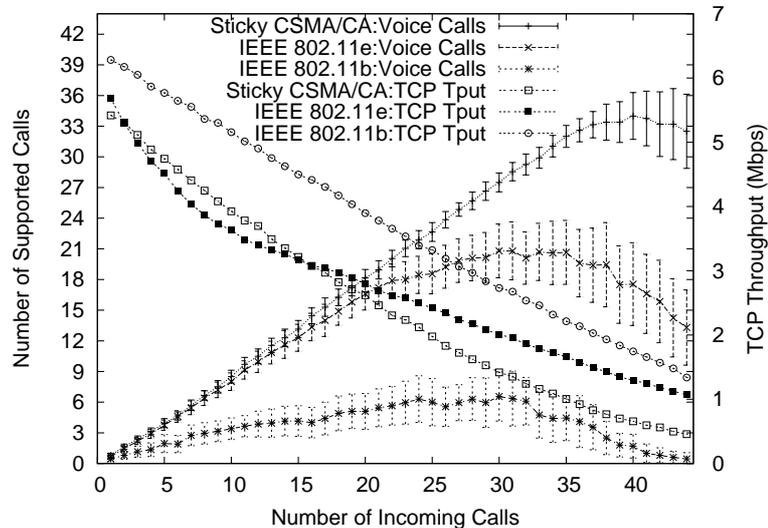


Figure 4.20: Voice call support for regular grid topology network with silence suppression enabled.

ity of packet transmissions over talk spurt durations, thereby reducing medium contention and packet loss due to collisions.

4.6 Summary

In this chapter, we have introduced Sticky CSMA/CA, a MAC framework that can exploit regularity in real-time traffic to achieve implicit synchronization over wireless mesh networks, thereby enabling TDM-like QoS for such applications. For the mix of voice and data traffic considered in our numerical results, Sticky CSMA/CA locks onto a periodic transmission schedule that matches the period of the VoIP flows. This periodicity in transmission enables each node to

accurately predict the real-time traffic generated by nodes that are capable of interfering with it, thereby drastically reducing the contention seen by real-time flows. Thus, Sticky CSMA/CA mechanism provides quasi-deterministic QoS to a real-time flow once it is established, while its slot reconfiguration mechanism enables robust recovery from losses due to unexpected interference patterns resulting from hidden terminals. Through analysis and simulations, we have shown that Sticky CSMA/CA significantly increases the total voice call carrying capacity of mesh networks as compared to the conventional CSMA/CA based schemes such as IEEE 802.11b/e. While we impose a strict priority for real-time flows over delay-insensitive data traffic, the new contention mechanisms introduced for data packet transmissions ensure that the data traffic effectively fills the bandwidth left over after serving the VoIP flows. Our simulation results for different scenarios show that the total voice call carrying capacity of Sticky CSMA/CA is much higher than that of IEEE 802.11b/e, while its data throughput remains comparable to that of IEEE 802.11b/e.

It is worth pointing out that, while IEEE 802.11e is intended to provide QoS to real-time flows by giving them priority over delay-insensitive traffic, this service differentiation has little effect when a number of real-time flows of equal priority contend for the medium. In particular, our performance evaluations over the grid topology (characterized by high interference and contention for the medium)

demonstrate that IEEE 802.11e yields much lower voice call carrying capacity and higher delay/delay jitter than Sticky CSMA/CA. In contrast, Sticky CSMA/CA achieves fine-grained resource allocation among nodes that have real-time packets to transmit, while enforcing strict service differentiation among real-time and delay-insensitive traffic.

While our performance evaluation was restricted to VoIP flows that naturally exhibit periodicity that can be exploited by Sticky CSMA/CA, the promising results that we obtain motivate research on extending these ideas to real-time flows that may not be naturally periodic. In particular, our ongoing research focuses on efficient support of variable bit rate (VBR) traffic such as video using Sticky CSMA/CA. We propose to impose a basic unit of resource allocation in the network as a periodic connection with packets of a fixed size, thus satisfying the basic requirements for Sticky CSMA/CA: implicit synchronization and prediction of network activity via Carrier Sense Tables. Nodes initiate and terminate such connections based on their bandwidth needs, analogous to the arrival and departure of voice connections in the scenario considered in this chapter. The specific criteria for opening and closing connections depends on the applications being supported. For VBR video, for example, it may be possible to use an approach similar to Renegotiated CBR, which was proposed in [107] for supporting video over ATM networks. This approach is based on splitting a VBR flow into

a base CBR stream with minimal bandwidth, and renegotiating for additional bandwidth in terms of additional CBR streams whenever required. The buffer occupancy levels at the application layer can be used to deduce the need for additional bandwidth and to estimate the bandwidth required. Significant cross-layer interaction is required to be able to provide a deterministic QoS. This interaction includes admission control to decide whether or not a new VBR flow should be admitted while ensuring that bandwidth renegotiations for existing flows are successful, end-to-end signaling for flow establishment and bandwidth renegotiations, and buffering at the application layer along with an intelligent algorithm to determine if bandwidth renegotiation is required for an existing VBR flow. The efficacy of such techniques for supporting VBR video over wireless mesh networks is the subject of our current research.

More broadly, Sticky CSMA/CA could potentially be used as a fundamental building block of a cross-layer architecture to provide comprehensive QoS solutions for wireless mesh networks. Specifically, Sticky CSMA/CA can be exploited by intelligent routing, transport and application layer schemes to provide deterministic QoS guarantees to applications with different QoS requirements. For example, a wireless mesh routing protocol can leverage the medium utilization information maintained in the CS Tables of nodes to search for QoS capable routes that can help in efficient resource utilization. Admission control algorithms can utilize the

CS Table maintained at each node to obtain an accurate estimate of the available resources in the network, to be able to decide whether a new flow should be admitted, and to provide deterministic QoS guarantees to existing flows. In order to enforce the desired resource allocation among different traffic classes and to prevent starvation, thresholds on maximum allocated bandwidth for each traffic class can be decided by the network designer and can be enforced through admission control schemes. By enabling local (per-hop) resource reservations, supporting service differentiation, and providing an estimate of available resources in the network, Sticky CSMA/CA supports QoS constrained routing and admission control, which are essential components of any QoS architecture.

Chapter 5

Shaping Throughput Profiles in Multihop Wireless Networks

5.1 Introduction

In this chapter, we study resource sharing strategies for partitioning the transport capacity of static multihop wireless networks. The throughput performance for these networks is determined by how the transport capacity is divided up among competing flows. Consider a static multihop wireless network of N identical nodes with a fixed link rate and spatially uniform traffic where each node randomly chooses a destination. For such a network, Gupta and Kumar [1] showed that the achievable per-flow throughput diminishes approximately as $O(\frac{1}{\sqrt{N}})$. Intuitively, this negative scalability result follows from the fact that for a network deployment area A , the total transport capacity only increases as $O(\sqrt{AN})$, whereas the number of hops to the destination and hence the amount of network resources required

per-connection scales up as $O(\sqrt{N})$. However, this per-flow throughput analysis implicitly assumes that the resource allocation to each connection is throughput-fair across the network. Therefore, connections that traverse a larger number of hops (longer connections) consume significantly more network resources than those traversing fewer hops (shorter connections) to achieve the same end-to-end throughput.

A natural question that arises from the preceding observation is what if we sidestep assumption of throughput-fair allocation by intelligently biasing network resource allocation against resource-intensive longer connections? Shorter connections could then achieve a significantly higher throughput at the cost of a controlled performance degradation for the longer connections. In this chapter, we introduce a rich class of resource allocation strategies that embody this design philosophy. Our goal is to provide the network designer with a flexible means of trading off throughput performance seen by long and short connections in a manner consistent with network design goals.

Resource sharing strategies from prior literature that fall within our framework are proportional fairness [108], and the more general (p, β) -proportional fairness [109]. However, we show that it is possible to design the resource allocation strategy to choose from among a much larger class of flow throughput profiles (i.e., flow throughput versus the number of hops the flow traverses) while main-

taining efficient network utilization. For example, it is possible to significantly improve the performance of shorter connections beyond that attained by proportional fairness, while minimally degrading the performance of longer connections. Alternatively, we can improve throughput for both short and long connections relative to proportional fairness, at the expense of slightly lower throughput for flows traversing a moderate number of hops.

While it is intuitive that taking away a small amount of resources from long connections can significantly improve performance for short connections, naively biasing against long connections can lead to poor performance. In a large network with spatially uniform traffic, a large fraction of total connections are long, so that there are not enough short connections to take advantage of the transport capacity that is released if we strongly bias against long connections. We draw upon this key observation to propose a new class of resource allocation strategies called *mixed-bias* strategies, in which a fraction of the total available network resource is allocated via an allocation strategy that is strongly biased against long connections, with the remainder allocated in an unbiased (e.g., max-min fair) or mildly biased (e.g., proportional fair) manner. Mixed-bias strategies enable the desired flexibility in shaping throughput profiles while maintaining efficient network utilization.

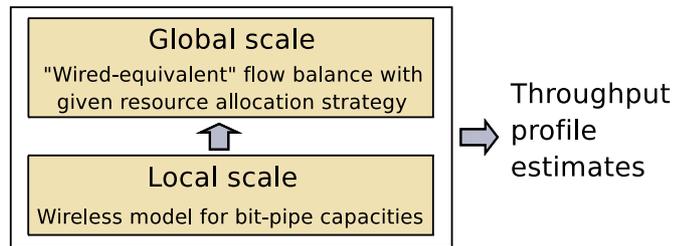


Figure 5.1: Two-scale model for resource allocation over large multihop wireless networks.

Our second contribution is an analytical model that provides quick estimates of the throughput profile achieved by a given resource allocation strategy. This model serves as a tool for exploring the rich design space of mixed bias strategies in shaping throughput profiles that are consistent with network performance goals. The foundation for our analytical framework is a novel two-scale model for large multihop wireless networks (see Fig. 5.1). At the *global* scale, the decay of transmit power with distance and the broadcast nature of wireless medium (which limits spatial reuse), necessitate the use of short distance hops. Therefore, wireless plays a fundamental role only in influencing the network topology. At this scale, elementary flow calculations that treat node bandwidth shares (or data transfer capacity) as bit pipes suffice to analyze the impact of various resource allocation strategies. For instance, for throughput-fair scheduling at the nodes, Gupta and Kumar’s negative scalability results follow from such flow analysis, and so do results for different biased resource allocation strate-

gies. The effective node data transfer capacities used in the preceding “wired-equivalent” global scale calculations are obtained from a detailed analysis of the wireless medium at the *local* scale, which takes into account medium access control, transceiver capabilities, and the radio propagation characteristics. While our approach is based on flow balance at a “typical” node in a large network, it is also refined to account for the effect of finite network size and spatial traffic patterns.

We verify the design prescriptions obtained from our analytical model for an example application over 802.11-based multihop networks. The simulations of IEEE 802.11-based multihop networks incorporate the effects of complex node interactions at different layers of the wireless network stack that are not explicitly modeled in our analysis. The performance trends from the simulations demonstrate that the two-scale model, despite its simplicity, captures the fundamental design tradeoffs.

For practical protocol translations of our resource biasing strategies, we map the throughput profiles attained via our resource allocation to a set of concave utility functions consistent with a network utility maximization (NUM) framework (further discussed in the Section 5.2). The extensive work on protocol designs for NUM over multihop wireless networks can therefore be applied to obtain protocols that can implement mixed bias strategies. Interestingly, we find that our

resource allocation strategies are not only consistent with the NUM framework, but they also extend the class of utility functions generally considered for resource allocation [109].

We now summarize our main contributions:

- We propose *mixed-bias* strategies that blend strongly biased resource allocation with fairer allocation strategies. We show that mixed-bias strategies achieve a superior flow throughput profile and efficient network utilization that cannot otherwise be attained by the individual strategies in the mixture.
- We present a two-scale analytical model that provides insight into the impact of a particular resource allocation strategy on network performance, in a manner that captures the effect of finite network size and spatial traffic patterns. Our two-scale model represents a simple but powerful approach to design and performance analysis of large scale multihop wireless networks. We present packet-level simulations for an example application to verify the performance predictions from our analytical model.
- We show that our resource allocation strategies are not only consistent with a NUM framework, but they also extend the commonly considered class of utility functions in the literature [109]. This opens the way to leverage extensive prior research on NUM-based protocols for future work on translating our resource biasing approach to implementable network protocols.

The rest of the chapter is organized as follows. Section 5.2 discusses the related work. Section 5.3 presents our resource allocation framework and describes mixed-bias allocation policies. In Section 5.4, we present an example application of our resource biasing framework and present a simulation evaluation. In Section 5.5, we establish the connection between our resource biasing strategies and a NUM framework, which shows that our resource allocation policies are amenable to distributed protocol implementations. We conclude with a summary of our findings in Section 5.6.

5.2 Related Work

Most prior work on resource allocation in communication networks is based on a network utility maximization framework [108]: a user (i.e., flow) i attaining rate x_i obtains utility $U(x_i)$, where $U(\cdot)$ is a strictly concave, non-decreasing, continuously differentiable function, and the goal of resource allocation is to maximize the sum of utilities, subject to network resource constraints. Consider a network with a set L of resources or logical links with capacities $R = (r_l, l \in L)$, and a set of flows F . The basic NUM problem formulation for a wired network is as follows [108]:

$$\max \sum_{i \in F} U(x_i) \tag{5.1}$$

subject to

$$AX \leq R \quad \text{over} \quad X \geq 0 \tag{5.2}$$

where $X = (x_i, i \in F)$ is the vector of user (flow) rates, $A = (A_{li}, l \in L, i \in F)$ is a 0 – 1 routing matrix where $A_{li} = 1$ if flow i uses link l , $A_{li} = 0$ otherwise. The above formulation is a convex optimization problem, given the concavity of the objective function (5.1) and the compact feasible region (5.2). In addition to convexity, typical approaches to solving NUM problems identify and exploit a decomposition structure embedded in the problem to design decentralized algorithms that converge to the global optimum. The solution algorithms are cast into network protocols that can span multiple layers of the network stack.

A utility function can be interpreted as a knob to trade off the network throughput performance versus fairness among competing flows. Reference [109] proposes the following class of utility functions which have been commonly considered in the subsequent work on NUM.

$$U(x, \beta) = \begin{cases} (1 - \beta)^{-1} x^{1-\beta} & \beta \neq 1 \\ \log(x) & \beta = 1, \end{cases} \tag{5.3}$$

where $x \geq 0$ is flow rate and parameter $\beta > 0$ captures the degree of fairness, where a higher β implies a more concave $U(x)$, and hence a fairer allocation. $U(x, \beta)$ maps to sum throughput maximization for $\beta \rightarrow 0$, proportional fairness [108] for

$\beta = 1$, potential delay minimization [110] for $\beta = 2$, and approaches max-min fairness [111] for $\beta \rightarrow \infty$.

It is worth pointing out that, unlike the above-mentioned NUM framework, our starting point is quite different, in that we explicitly assume that the resource allocation for a flow is proportional to a specific decreasing function of the number of hops traversed by a flow, and then compute the proportionality constant via flow balance. However, we show in Section 5.5 that our biasing strategies are consistent with a NUM framework in the sense that throughput profiles achieved via our biased allocation policies map to specific concave utility functions under a few mild assumptions over large multihop wireless networks. This connection indicates that the recent work on NUM algorithms over multihop wireless networks [112–120] can be exploited for protocol translations of our resource biasing strategies.

Following the seminal work of Kelly [108], there has been extensive research on NUM in the context of resource allocation problems in wireline networks, and more recently, wireless networks. Comprehensive surveys of work on the applications of the NUM framework can be found in [2, 3, 121]. For multihop wireless networks, [119, 120] propose congestion control algorithms based on NUM formulations. More recent papers [112–117] consider joint congestion control and scheduling problem to design cross-layer resource allocation algorithms using the NUM framework. Most of these papers basically employ combinations of ideas

from distributed wireline congestion control (e.g., [108] and follow up work), and queue-lengths based wireless scheduling or random access MAC (many of which are inspired from [122] and its extensions). Our work is complementary to this body of work. By connecting our mixed bias strategies to the NUM framework, we show that the network designer actually can work with a much broader class of utility functions than considered in prior work. On the other hand, putting our resource allocation strategies within the context of utility maximization implies that prior work on NUM-based protocols (e.g., [112–117]) can be leveraged to translate our framework to implementable protocols.

Other relevant prior work includes a discussion of the tradeoff between network capacity and fairness for heterogeneous multihop wireless networks [123], which concludes with advocacy of proportional fairness as a reasonable compromise between rate efficiency and fairness, and also observes that max-min fairness yields poor performance because all flows obtain essentially the same rate as the minimum rate flow. Reference [124] presents a realization of proportional fairness based on end-to-end rate control for a connection based on the number of hops it traverses.

5.3 Resource Allocation Framework

We now describe our model for the resource allocation framework. We first discuss the global scale analysis, which essentially involves flow balance in a wired-equivalent network whose capacities would be set by a local scale analysis.

For a uniform traffic model, each node generates a persistent flow, choosing a destination at random from among all other nodes. Our framework also accommodates spatially localized traffic distributions in which the probability of choosing a destination depends on the number of hops from the source. The biased resource allocation problem is to assign bandwidth to each flow as a function of its resource requirements (e.g., number of hops it traverses) such that node data transfer capacity constraints are honored. Since it is intractable to identify bottleneck links for randomized traffic patterns, we apply an averaged analysis at the global scale to do flow balance for a “typical” node of the following form:

$$C \geq \sum_{h=1}^{h_{max}} n(h)c(h), \quad (5.4)$$

where C is the available data transfer capacity or bandwidth share of a node (in bits per-second), $n(h)$ is the average number of h -hop connections traversing the node, $c(h)$ is the effective data rate assigned to h -hop flows, and h_{max} is the maximum number of hops. Note that $c(h) \leq s$, where s is the maximum stable data rate of a link. Both C and s are determined via a detailed local scale

analysis of wireless medium access control and physical layer properties: s can be considered as the data transfer capacity of a node's wireless contention region, whereas C is a node's effective share of s . For an N -node network, $h_{max} = O(N)$ for a linear topology, and $h_{max} = O(\sqrt{N})$ for a two-dimensional grid.

For a uniform traffic model, a typical node is more likely to see a longer flow (since such a flow traverses more hops), so that $n(h)$ increases with h . It therefore makes sense to make $c(h)$ a decreasing function of h , in order to ensure that the sum in (5.4) does not exceed C even as the network size (and hence h_{max}) becomes large. A throughput profile refers to the dependence of $c(h)$ with h . Our goal is to explore the design space of throughput profiles that are feasible, in terms of meeting the link balance condition (5.4).

For simplicity of exposition, we first illustrate our framework with some quick global scale computations for a one-dimensional network in Section 5.3.1. We then provide a more comprehensive analysis for a two-dimensional network which models the IEEE 802.11 network to be simulated.

5.3.1 Computations for a One-Dimensional Network

Consider a linear network with N nodes (arranged as a ring to avoid edge effects: see Fig. 5.2) and $C = 1$. For uniform traffic, the probability of a connection initiated by any given node traversing h hops is $p(h) = \frac{1}{N-1}$, $h = 1, \dots, N - 1$,

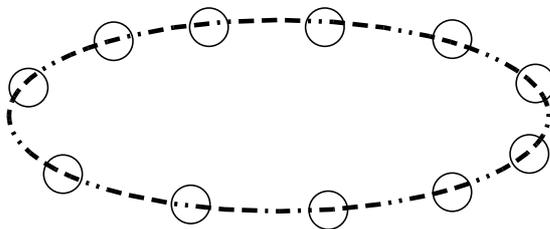


Figure 5.2: Linear topology.

since any of the other $N - 1$ nodes in the network can be chosen with equal probability. The average number of h -hop connections traversing a node is given by $n(h) = \frac{Np(h)h}{N} = hp(h)$, because there are on average $Np(h)$ flows that traverse h hops in a network of N nodes. Thus, flow balance at a typical node becomes (assuming full utilization):

$$\sum_{h=1}^{N-1} hp(h)c(h) = 1 \quad (5.5)$$

For a throughput-fair allocation, the bandwidth allocated to an h -hop connection is independent of h , so that $c(h) = \lambda$. We show that the preceding throughput profile maps to max-min fairness in Section 5.5. Solving (5.5), we obtain $c(h) = \lambda \approx \frac{2}{N}$. The sum throughput $T_{sum} = N \sum_h c(h)p(h)$, summed over all connections, is approximately 2.

Now consider the case when the bandwidth assigned to a connection is inversely proportional to the number of hops it traverses; i.e., $c(h) = \frac{\lambda}{h}$. Here λ denotes the maximum allowed source application data rate (in bits per-second). We show that this allocation approximates proportional fairness in Section 5.5.

Plugging this into (5.5), we obtain that $\lambda \approx 1$, so that $c(h) \approx \frac{1}{h}$. The sum throughput $T_{sum} \approx \log N$. Thus, we have significantly improved both the sum throughput and the performance of shorter connections, with only a factor of two loss in the throughput obtained by the longest connections. One can show using typical link flow analysis that the throughput of a h -hop connection is approximately $\frac{1}{h}$, with the sum throughput of $\log N$, and $\lambda = O(1)$. Comparing with the max-min fair allocation, we realize that biasing against the long connections via proportional fairness results in a win-win strategy: it leads to far better performance for the shorter connections, while reducing the throughput attained by the long connections only by roughly a factor of two. The cumulative effect of this is the improvement of the total throughput by a factor of $\frac{\log N}{2}$ compared to the max-min fair allocation.

The pitfalls of naïve biasing: Given the improved performance tradeoff offered by proportional fairness, it is natural to ask if we can get still larger improvements by biasing even more severely against the long connections. Consider a more severely biased allocation $c(h) = \frac{\lambda}{h^2}$. Substituting into (5.5), it can be shown that there is indeed a gain in sum throughput, but we require that $\lambda = O(\frac{N}{\log(N)})$. That is, the maximum flow throughput must scale up with network size, which is clearly not compatible with the finite capacity of one. Basically, there are not enough short flows to take advantage of the transport capacity released by such a

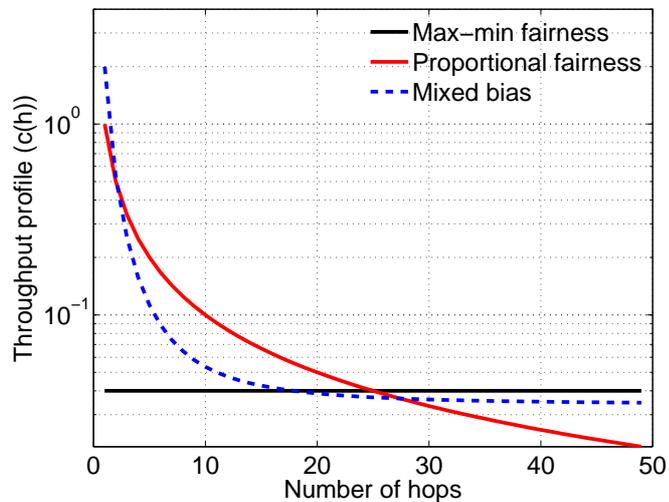


Figure 5.3: Throughput profiles for mixed biasing over 50 node linear (ring) topology.

strong bias. If we now impose the constraint that λ cannot scale up with N , we find (details are provided for a two-dimensional model later) that the network is poorly utilized.

The promise of mixed biasing: The problem with the preceding biasing strategy was that the short connections would have to send at too high a rate (more than the maximum stable link rate s) to take advantage of the bias. However, if we only allocate a small fraction of the available capacity to this biasing strategy, then we can eliminate this problem. For example, consider a *mixed bias* strategy of the form $c(h) = \lambda_1 + \frac{\lambda_2}{h^2}$; this mixes max-min fairness with a bias stronger than proportional fairness. Let us assume that we use a proportion α_1 of the node data transfer capacity to implement max-min fairness, and a proportion $\alpha_2 = 1 - \alpha_1$ to

implement the stronger bias such that flow balance decomposes into two separate equations:

$$\sum_{h=1}^{N-1} hp(h)\lambda_1 = \alpha_1, \text{ and } \sum_{h=1}^{N-1} hp(h)\frac{\lambda_2}{h^2} = \alpha_2.$$

Now, by choosing α_2 to scale down with N , we can prevent λ_2 from scaling up with N . Specifically, we can set $\alpha_2 = a\frac{\log N}{N}$, which will result in $\lambda_2 \approx a$, which no longer scales up with N . The resulting throughput profile $c(h)$ is a convex combination of the throughput profiles from the two different strategies:

$$c(h) \approx \frac{a}{h^2} + \left(1 - a\frac{\log N}{N}\right) \frac{2}{N}. \quad (5.6)$$

Here, we constrain $a \leq \frac{s-2/N}{1-2\log(N)/N^2}$ so that the maximum throughput does not exceed the stable link rate s . We note that (5.6) provides a very different throughput profile than either max-min fairness or proportional fairness. In particular, the throughput seen by long connections is almost as good as that of max-min fairness (and hence is a factor of two better than that obtained from proportional fairness), while the throughput seen by short connections does not scale down as a function of N (and hence is better than max-min fairness). Fig. 5.3 illustrates example throughput profiles for a 50-node linear network with link data rate $s = 2$ and node bandwidth share $C = 1$. In short, mixed bias strategies open up a large design space not covered by existing resource allocation strategies. We shall explore this in more detail in our discussion of two-dimensional networks.

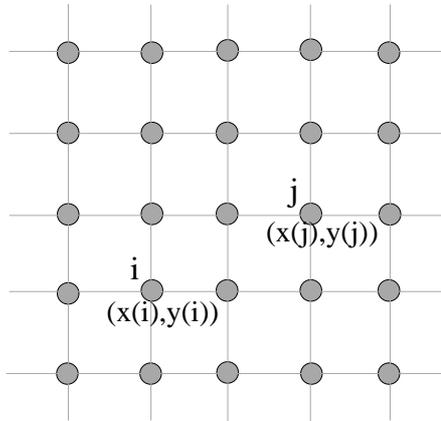


Figure 5.4: Regular grid topology.

5.3.2 Two-Dimensional Networks

We now consider two-dimensional N node regular grid topology networks where the inter-node distance is such that direct communication is only possible between immediate neighbors on the grid. The hop distance $h_{(i,j)}$ between two nodes i and j with coordinates (x_i, y_i) and (x_j, y_j) , respectively, is given by $h_{(i,j)} = |x_i - x_j| + |y_i - y_j|$, which equals the number of hops along the shortest path between the nodes (see Fig. 5.4).

We consider a power-law traffic model [125] where the probability that a node i will communicate with a node j , h hops away, is given by

$$p_{i,j} = \frac{q}{h^k}, \quad (5.7)$$

where $0 < q < 1$. Note that $p_{i,j} = p_{j,i} = p(h)$. Thus, the probability $p_H(h)$ that a pair of communicating nodes are h hops apart is given by

$$p_H(h) = \frac{m(h)}{h^k \sum_{r=1}^{h_{max}} \frac{m(r)}{r^k}}, \quad (5.8)$$

where $m(h)$ is the total possible number of h -hop flows in the network. For $k > 0$, this spatial traffic distribution models a spatially localized traffic pattern with degree of localization increasing with k . We first consider uniformly distributed traffic ($k = 0$) to study the effect of different resource biasing policies.

We focus on allocation of the available data transfer capacity C (in bits per-second) of a typical network node using link flow analysis to obtain the expected per-connection throughput. The underlying medium access control (MAC) layer (e.g., IEEE 802.11) implements max-min fair bandwidth allocation among contending nodes. The bias weight function $w(h)$ is defined as $w(h) = \frac{1}{h^b}$, where the bias exponent $b \geq 0$ determines the degree of bias. Let $n(h)$ denote the average number of h -hop flows passing through or initiated by a typical network node.

We have

$$C \geq \sum_{h=1}^{h_{max}} n(h) \frac{\lambda}{h^b}. \quad (5.9)$$

Thus, the required λ for maximum utilization of the available capacity is given by $\lambda = C / \sum_{h=1}^{h_{max}} \frac{n(h)}{h^b}$.

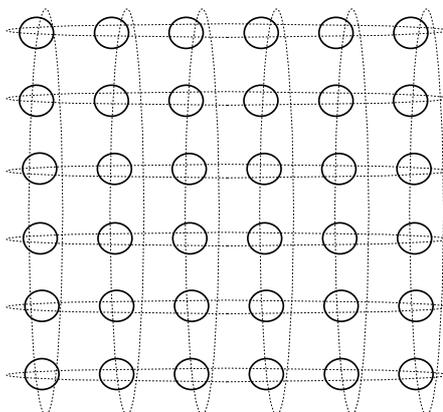


Figure 5.5: Two-dimensional torus topology.

We first consider a symmetric, regular grid (i.e., a 2-d torus, as shown in Fig. 5.5) topology to avoid edge effects, and focus on how the required λ scales with increasing network size as a function of the traffic model and the bias weight function. For a symmetric regular grid topology, $p_H(h) = 1/(h^{(k-1)} \sum_{r=1}^{h_{max}} \frac{1}{r^{(k-1)}})$. This is because the number of nodes located h hop away from any given node is $4h$, which yields the total possible number of h hop flows as $4Nh$. This is plugged in (5.8) to obtain $p_H(h)$. With the average number of h hop connections traversing a node given by $n(h) = hp_H(h)$, we have from (5.9),

$$\lambda = \frac{C}{\sum_{h=1}^{h_{max}} hp_H(h)/h^b} = \frac{C \sum_{h=1}^{h_{max}} 1/h^{(k-1)}}{\sum_{h=1}^{h_{max}} 1/h^{(k+b-2)}}. \quad (5.10)$$

For uniform traffic distribution ($k = 0$), we infer that with unbiased resource allocation ($b = 0$), $\lambda = O(1/\sqrt{N})$; for proportional fairness ($b = 1$), we have $\lambda = O(1)$; and for $b = 2$, $\lambda = O(\sqrt{N})$. Further calculations show that the required λ grows

with the network size N for $b > 1$. However, since λ cannot exceed the maximum stable data rate s , this means that after the network size increases to a certain limit for $b > 1$, a node cannot fully utilize its fair share of bandwidth because the stronger bias prevents long flows from using the leftover capacity from the short flows. As in the linear network case, we infer that with proportional fair resource allocation, the per-flow throughput will be higher than the unbiased resource allocation for any N since the required λ is $O(1)$. Note that the preceding analysis also applies to infinite regular grid topologies where edge effects can be neglected.

Spatially localized traffic distribution: We now consider the effect of spatially localized traffic ($k > 0$). From (5.9) and (5.10), we find that for $k > 3$ in $p_H(h)$ defined in (5.8), the per-flow throughput stays constant, i.e., $O(1)$, rather than diminishing with increasing network size for unbiased resource allocation ($b = 0$). Even for traffic distributions with $2 < k \leq 3$, biased resource allocation with $b > 3 - k$ can ensure that per-flow throughput does not diminish. The observations about spatially localized traffic patterns and unbiased resource allocation are consistent with insights from [1, 126].

Finite, regular grid topologies: We now consider finite, regular grid topologies where the contention at the network edges is less than the core because there are fewer contending nodes at the edges; also, with a uniform traffic model where every node randomly picks a destination among other network nodes, the traffic volume

increases towards the core of the network. Therefore, the network core acts as a bottleneck for network flows in most cases. Hence, we focus on the network-core to avoid offsets from the edge effects. We expect that the flow throughput results obtained from analysis of the network core are pessimistic estimates of the actual flow throughputs, especially for smaller networks (e.g., $N < 50$) where the edge effect dominates. Based on these observations, we modify (5.9) as follows:

$$C \geq \sum_{h=1}^{h_{max}} n_c(h) \frac{\lambda}{h^b}, \quad (5.11)$$

where $0 < \lambda \leq s$ and $n_c(h)$ denotes the average number of h -hop flows passing through or initiated by a typical network-core node. We evaluate $n_c(h)$ by averaging the values obtained from MatLab simulations of traffic instances over the network for 5000 seed values. Thus, we can calculate the maximum λ that satisfies (5.11) for a given bias exponent b . This λ value can be used to estimate the average per-flow throughput T of the network as a function of the network size, traffic distribution and the resource allocation policy as follows:

$$T = \sum_{h=1}^{h_{max}} p_H(h) \frac{\lambda}{h^b}. \quad (5.12)$$

Mixed Bias Resource Allocation Policy: The underlying idea is to allocate a portion of the total available capacity at a node via a strongly biased policy, and allocate the rest employing a fairer policy. Assume that, of the total available node capacity C , we allocate αC via a resource allocation policy determined by

weight function $w_1(h) = \frac{1}{h^{b_1}}$, and allocate the remaining bandwidth via weight function $w_2(h) = \frac{1}{h^{b_2}}$. We calculate the maximum λ_1 and λ_2 ($0 < \lambda_1, \lambda_2 \leq s$) that satisfy the following inequalities:

$$\alpha C \geq \sum_{h=1}^{h_{max}} n_c(h) \frac{\lambda_1}{h^{b_1}}, \quad (5.13)$$

$$(1 - \alpha)C \geq \sum_{h=1}^{h_{max}} n_c(h) \frac{\lambda_2}{h^{b_2}}, \quad \lambda_1 + \lambda_2 \leq s \quad (5.14)$$

The average per-flow throughput in this case is given by:

$$T = \sum_{h=1}^{h_{max}} p_H(h) \left(\frac{\lambda_1}{h^{b_1}} + \frac{\lambda_2}{h^{b_2}} \right). \quad (5.15)$$

The parameters α , b_1 and b_2 allow us to span a large design space of resource allocation strategies, trading off throughput performance versus fairness. Note that an alternative formulation to (5.14) that ensures that any left-over capacity from the fraction of bandwidth allocated to strongly biased allocation in (5.13) is used up by the fairer allocation is given by $C \geq \sum_{h=1}^{h_{max}} n_c(h) \left(\frac{\lambda_1}{h^{b_1}} + \frac{\lambda_2}{h^{b_2}} \right)$, where $\lambda_1 + \lambda_2 \leq s$. Since our intent is to first understand the effect of different parameters on the throughput profiles, we focus on (5.14) for our evaluations. We provide insight in to the effect of these parameters in the next section for an example application of our framework over IEEE 802.11 multihop networks.

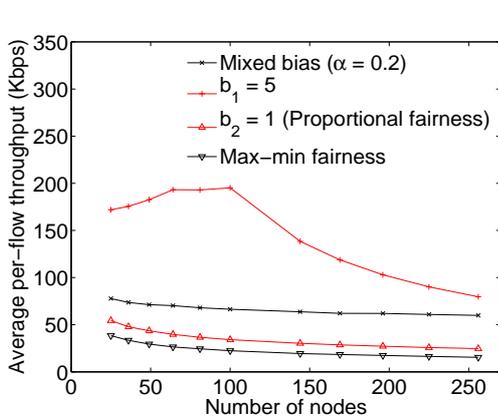


Figure 5.6: Average per-flow throughput for mixed bias with $b_1=5$, $b_2=1$, and $\alpha = 0.2$ from analysis.

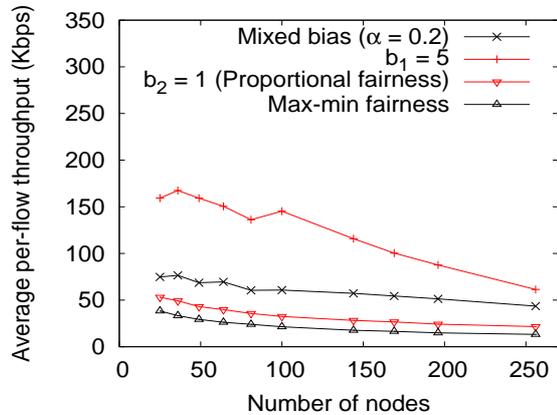


Figure 5.7: Average per-flow throughput for mixed bias with $b_1 = 5$, $b_2 = 1$, and $\alpha = 0.2$ from simulations.

5.4 Application to an IEEE 802.11 Multihop Network

We now apply our resource allocation framework to an IEEE 802.11 multihop network to illustrate how to determine global-scale parameters from local-scale analysis, which takes into account physical layer technology and medium access control. We then employ simulations with IEEE 802.11b parameters to determine how well the design prescriptions obtained by our analysis in Section 5.3 work.

For our local-scale analysis, Bianchi's saturation throughput analysis [95] is used to determine the data transfer capacity s of a contention region in an IEEE 802.11 network. Based on the observation that the IEEE 802.11 MAC protocol

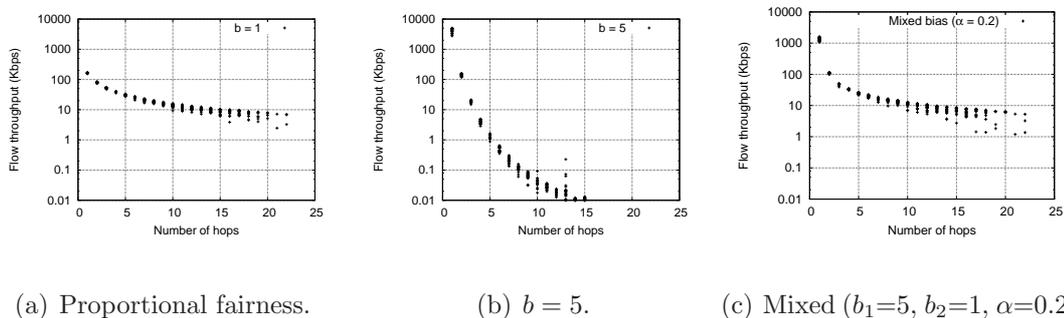


Figure 5.8: Scatterplots of flow throughputs for a 144-node network for mixed bias allocations, proportional fairness and severe bias ($b = 5$).

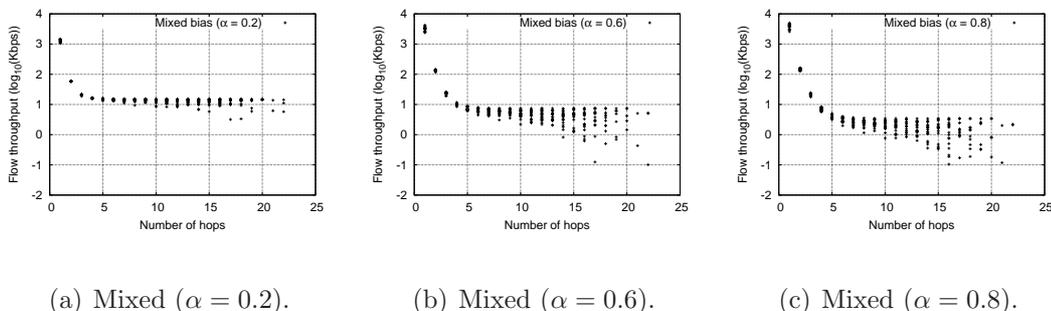


Figure 5.9: Scatterplots of flow throughputs for a 144-node network for mixed bias allocation policies ($b_1 = 5, b_2 = 0$).

implements an approximation to max-min fair bandwidth allocation among contending neighbors [123, 127], a typical node's share of single hop capacity can be estimated as $C = s/n_c$, where n_c is the average number of contending nodes in a contention region at the network core. This value of C is now used in (5.11) in the global-scale analysis to estimate λ . This is then plugged into (5.12) to obtain the average per-flow throughput. Fig. 5.6 illustrates the predicted average per-flow throughput for an example mixed bias allocation with $b_1 = 5, b_2 = 1$ (propor-

tional fairness) and $\alpha = 0.2$, for the individual allocation policies in the mixture and also for max-min fair allocation. We observe that the mixed bias allocation has a higher average flow throughput as compared to proportional fair and max-min fair allocations. Evaluation for different α values shows that increasing α to a higher value emphasizes strong bias whereas decreasing α to a smaller value such as 0.2 sways the throughput more towards the fairer allocation in the mixture.

We now employ simulations to obtain flow throughput profiles, and to evaluate the benefits of mixed bias strategies relative to the individual strategies in the mixture, as well as to proportional fairness. We use the QualNet Network Simulator [61] over regular grid topology networks with IEEE 802.11b MAC/PHY, in basic access mode (RTS-CTS turned off). We present our results with static routing in order to isolate the effects of routing overhead on throughput. The results are averaged over 20 different seeds. Our application model comprises CBR flows of packet size of 1000 bytes, with data rates determined analytically for a given biasing strategy, as in Section 5.3.2 and enforced using source data rate control. Destinations are chosen at random based on the uniform traffic model.

While we fix the offered rate based on the analytical model in order to compare different resource allocation strategies, as we show in the next section, embedding mixed-bias resource allocation strategies within a utility function framework can provide the flexibility needed to arrive at decentralized adaptive implementations

that allow nodes to dynamically tune their flow rates while implementing a biased resource allocation strategy. This can be used to alleviate congestion, to exploit spatial reuse opportunities, and more broadly, to react to specific network topologies and traffic patterns. In fact, embedding our biasing strategies on a NUM framework allows us to leverage protocol translations of jointly optimal cross-layer congestion control and scheduling algorithms, rather than restricting ourselves to flow rate control over a given MAC, as in this evaluation. Our aim in this section is to illustrate the promise of mixed biasing strategies even in this restrictive setting.

Fig. 5.7 illustrates the per-flow throughput performance for the example mixed bias allocation considered in Section 5.3.2 ($b_1 = 5, b_2 = 1, \alpha = 0.2$) and compares it with that of the single bias policies in the mixture and also max-min fair allocation. The results are consistent with the analytical prediction that the mixed bias allocation provides higher average flow throughput than max-min fair or proportional fair allocations. However, the average flow throughput values are lower, and decay faster with network size, than the analytical predictions shown in Fig. 5.6. This is because flows incur higher packet loss probabilities as they traverse a larger number of hops in a larger network, which results in an increase in the amount of wasted network resources. Another factor causing performance degradation that is not explicitly modeled in the analysis is the waste of medium access time due to

the interaction between the IEEE 802.11 backoff mechanism and multihop packet forwarding [126].

Fig. 5.8 presents the scatterplots of individual flow throughputs as functions of the number of hops for a 144-node network over all simulation seeds. We consider a mixed bias strategy, mixing a strong bias with proportional fairness. The strategy leads to significantly higher throughput for shorter connections than proportional fairness, while incurring virtually no degradation for the longer connections. Other desirable mixed bias strategies that, for example, provide better performance for long connections than proportional fairness, can be obtained by mixing a strongly biased strategy with max-min fairness (e.g., Fig. 5.9(a)). Note that the strongly biased strategy shuts out the long connections, and is therefore not a feasible choice on its own. The choice of the capacity fraction α and the strength b of the bias are other parameters at the disposal of the system designer to craft a desirable throughput profile.

We now investigate the effect of different values of the capacity fraction α allocated via high bias. Consider the mixed bias allocation with $(b_1 = 5, b_2 = 0)$. Fig. 5.10 illustrates the throughput performance obtained from the simulations for the mixed bias policies for $\alpha = 0.2, 0.6, \text{ and } 0.8$. The figure also plots single bias policies in the mixture and proportional fair allocation for reference. Figs. 5.9(b) and 5.9(c) present the scatterplots of individual flow throughputs as functions

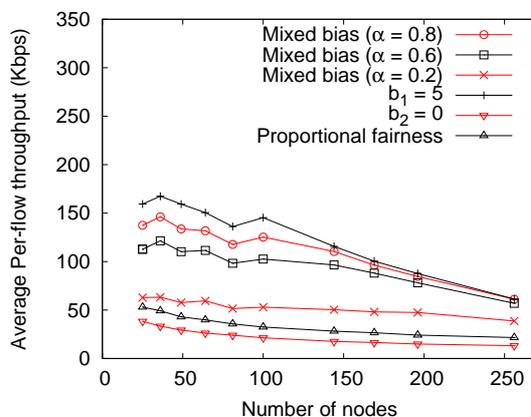


Figure 5.10: Average per-flow throughput for mixed bias with $b_1 = 5$, $b_2 = 0$, and $\alpha = 0.2, 0.6$, and 0.8 from simulations.

of the number of hops for a 144-node network over all simulation seeds, for the mixed bias allocations with capacity fraction $\alpha = 0.6$ and 0.8 (see Fig.5.9(a) for $\alpha = 0.2$). A comparison between Figs. 5.9(b) and 5.9(c) illustrates the poor network utilization with a higher bias and high capacity fraction α . We see that allocation of 80% of the node capacity via a strong bias (i.e., $\alpha = 0.8$, $b_1 = 5$) results in decrease of throughput achieved by long flows as compared to the case when $\alpha = 0.6$, but no apparent gain is observed in the throughput performance of the shorter flows. This demonstrates inefficient network utilization with a predominantly higher bias allocation for uniform traffic distribution. With different choices of capacity fraction α , it is possible to achieve a range of throughput profiles depending on the set of constituent biases in the mixture, the network size, and the traffic distribution.

5.5 Mapping to Network Utility Maximization

We now demonstrate that our resource biasing strategies are consistent with the network utility maximization (NUM) framework. Consider the convex optimization problem (5.1)–(5.2), for which a unique optimal solution exists [128]. From the Karush-Kuhn-Tucker conditions, the solution satisfies

$$U'(x_i) = \sum_{l:l \in i} \mu_l \quad \forall i \in F \quad (5.16)$$

where $\mu = (\mu_l, l \in L)$ is a vector of Lagrange multipliers or shadow prices for the links, $\{l : l \in i\}$ denotes the set of links traversed by flow i , and X in (5.2) and μ satisfy

$$AX \leq R, \quad \mu \geq 0, \quad \text{and} \quad \mu^T(AX - R) = 0. \quad (5.17)$$

Now consider (5.1)–(5.2) as the NUM problem for the global scale wired-equivalent model of a multihop wireless network (see Section 5.3). Under the assumptions of a large network where all bit-pipe capacities are approximately fully utilized (i.e., $AX \approx R$), and equal shadow price for each bit-pipe $\mu_l = \mu_c, l \in L$, we have from (5.16) that

$$U'(x_i) = h_i \mu_c, \quad \forall i \in F \quad (5.18)$$

where h_i is the number of hops traversed by flow i . First consider the resource biasing throughput profile $x_i(h_i) = \lambda h_i^{-1}$, where $\lambda > 0$ is calculated via flow

balance as described in Section 5.3. We now verify our claim that this throughput profile maps to a proportional fair allocation. We have $h_i = \frac{\lambda}{x_i}$, for which, we find that

$$U(x_i) = k \log(x_i), \quad (5.19)$$

where $k = \lambda\mu_c$ satisfies (5.18). This form of utility function $U(x_i)$ corresponds to proportional fairness. Note that the extension for weighted proportional fairness (i.e., the utility function corresponding to each $i \in F$ is scaled with a weight $w_i > 0$) is straightforward.

We next consider a biased allocation with $x_i(h_i) = \lambda h_i^{-b}$, where $\lambda, b > 0$. Hence, $h_i = \lambda x_i^{-\frac{1}{b}}$. We find that

$$U(x_i) = k \frac{x_i^{1-\beta}}{1-\beta}, \quad (5.20)$$

where $\beta = \frac{1}{b}$ and $k = \lambda\mu_c$, satisfies (5.18). Note that (5.20) is of the same form as the utility function for (p, β) -proportional fairness proposed in [109], with $\beta = \frac{1}{b}$.

There are a few important points to note here: first, most prior works that consider the class of utility functions proposed by [109] focus on $\beta \geq 1$, which maps to proportional fairness ($\beta = 1$), delay minimization ($\beta = 2$), and other utility functions with higher throughput fairness, approaching max-min fairness as $\beta \rightarrow \infty$. These functions map to our biased throughput profiles with bias exponent $1 > b > 0$. For instance, we can now show that flow throughputs

corresponding to max-min fair allocation are equal in this setting. Basically, we consider $x_i(h_i) = \lim_{b \rightarrow 0} \lambda h_i^{-b} = \lambda$, where parameter λ is the rate allocated to each flow, derived via flow balance.

Utility functions of the form (5.20) with $(0 < \beta < 1)$ map to the strongly biased cases $(b > 1)$. Recall that we showed in Section 5.3 that allocating all network resources via such policies can lead to poor network utilization for large networks. This insight served as a motivation for mixed bias resource allocation policies. It is instructive to realize that our assumption of near-full utilization of bit-pipe capacities for large networks in the derivation of (5.20) will hold only if R in (5.2) represents an appropriately chosen fraction of the actual bit-pipe capacities (e.g., $r_l = \alpha C \forall l$).

An equivalent NUM formulation for the mixed biasing policies over the wired-equivalent global model can now be developed as follows. Consider the following NUM problem

$$\max \sum_{i \in F} U_1(x_{i,1}) + U_2(x_{i,2}) \quad (5.21)$$

subject to

$$AX_1 \leq \alpha R, \quad X_1 \geq 0 \quad (5.22)$$

$$AX_2 \leq (1 - \alpha)R, \quad X_2 \geq 0 \quad (5.23)$$

where $U_1(\cdot)$ is the utility function corresponding to the strongly biased allocation strategy (e.g., (5.20) with $0 < \beta < 1$), $U_2(\cdot)$ corresponds to a fairer allocation such as max-min fairness or proportional fairness, $\alpha > 0$ represents the fraction of total resources allocated via the strongly biased allocation, and the flow rate vector $X = X_1 + X_2$. For example, mixing proportional fairness with a higher bias strategy with bias exponent $b = 5$ would lead to $U(x_i) = (\frac{(x_{i,1})^{0.8}}{0.8} + \log(x_{i,2}))$, The utility function could also be weighted as $w_i U(x_i)$ where $w_i > 0$ is the weight for flow i . The total rate for flow i is $x_i = x_{i,1} + x_{i,2}$. The form of the mixed biasing optimization problem is such that it can be decomposed into two parallel NUM problems in a straightforward manner, one with objective function $\sum_{i \in F} U_1(x_{i,1})$, under constraints (5.22), and the second with objective function $\sum_{i \in F} U_2(x_{i,2})$ and constraints (5.23).

5.6 Summary

The resource biasing framework presented in this chapter opens up a rich design space for sharing transport capacity in multihop wireless networks that goes beyond existing paradigms such as max-min fairness and proportional fairness by blending strongly biased and fairer allocations to get superior throughput profiles with high network utilization. Our two-scale model for multihop wireless networks

yields quick performance estimates for different biasing strategies, providing a tool that allows system designers to tune parameters so as to shape throughput profiles while maintaining network efficiency. For 802.11 networks with fixed link speeds, the global scale, wired-equivalent model provides predictions of throughput profiles that match trends obtained by simulation. Embedding our resource biasing strategies within a utility function framework shows that the extensive work on distributed protocol designs for network utility maximization (NUM) over wireless multihop networks can be leveraged for effective protocol translations that adapt to the network topology and traffic patterns. For instance, [112] addresses some practical design and implementation issues related to the joint congestion control and scheduling algorithm proposed in [129] and offers distributed protocols for NUM. A detailed study of modified versions of such protocols that implement mixed bias strategies is a topic for future research.

Chapter 6

Conclusions and Future Work

In the Introduction, we discussed that next generation wireless networks must be equipped to handle the demands of the rapidly growing wireless user base and the diverse applications. Of the many design aspects of wireless networking systems, medium access control (MAC) and resource allocation are among the most critical. Efficient medium access control that accounts for the physical characteristics of the wireless network links, transceiver capabilities, and the driving applications, along with flexible resource allocation, are key to enabling higher network capacity and resource utilization.

At a high level, the contributions of this dissertation can be classified into three broad categories: the first of these is the design of practical network architectures for multiGigabit wireless networking in the unlicensed 60 GHz millimeter (mm) wave band. Besides addressing some fundamental design bottlenecks in both indoor and outdoor mm wave networks and highlighting many research challenges

that still remain, our work serves to demonstrate the promise of this next frontier for high-speed wireless networking. The second direction is to rethink MAC protocols in 802.11 mesh networks, where we exploit application layer regularities and apply ideas of observation-based learning and memory to achieve time division multiplexing (TDM)-like resource sharing without network-wide synchronization. The third category is resource allocation in multihop wireless networks, where we propose a framework for investigating the effect of different resource allocation policies on network performance, and present strategies that can achieve a superior tradeoff between network throughput and fairness while ensuring high network utilization.

6.1 The Big Picture

A key idea instrumental to our approach to medium access control in wireless networks is application of learning and memory to attain implicit coordination in a decentralized fashion. We believe that our work on Sticky CSMA/CA and Memory-guided Directional MAC (MDMAC) just scratches the surface of the set of potential applications of the simple yet extremely powerful concept of *stigmergy* in addressing challenging access coordination and resource sharing problems in multihop wireless networks.

Our work on 60 GHz indoor wireless personal area networks (WPANs) and outdoor mesh networks illustrates the crucial role of cross-layer considerations in exploiting the abundant unlicensed spectrum in the 60 GHz band. Our 60 GHz MAC protocol designs are driven by modeling and analysis of the distinct traits of mm wave links given their peculiar propagation characteristics and the expected network operating environments. Sticky CSMA/CA, on the other hand, utilizes the carrier sensing capability of mesh nodes in the 2.4 and 5.8 GHz WiFi bands, and shapes or exploits application characteristics in the design of medium access control. The use of these features to achieve stigmergic coordination results in much better performance for a mix of voice over IP (VoIP) and data traffic relative to 802.11 WiFi MAC protocols.

Another key theme of our work on 60 GHz network architectures is the design of fully directional MAC protocols, i.e., protocols that do not assume any form of omnidirectional communication for control or coordination. In such settings, conventional approaches based on carrier sensing become ineffective because nodes cannot sense their surroundings in an omnidirectional fashion to derive feedback regarding transmission activity in their neighborhood, which makes transmit-receive coordination challenging among neighbor nodes. Moreover, in absence of omnidirectional sensing, neighbor discovery also becomes an important design constraint. Our directional networking proposals address these challenges.

Lastly, our simple two-scale analytical model for large multihop wireless networks (Chapter 5) is fundamental to design of resource allocation strategies that trade off performance seen by resource-intensive long connections and lightweight short connections. The key insight behind this model is that at a global scale, wireless communication plays a fundamental role only in influencing the network topology. That is, the broadcast nature of wireless medium and the decay of signal power with distance constrain us to use short distance hops. At this scale, elementary flow-balance over a “wired-equivalent” global-scale model suffices to study the impact of a given resource allocation strategy over a given spatial traffic pattern. The bit-pipe capacities used in the global-scale analysis are obtained via a detailed wireless model at the local scale.

6.2 Future work

There are a host of both broad and specific research issues for future research in protocol design for mm wave networks and WiFi mesh networks, and resource allocation over wireless networks, of which we mention but a few.

In the context of 60 GHz indoor WPANs, an important consideration in cross-layer design is the impact of antenna design and node form factor. We envision the use of circuit board antennas for consumer electronics devices, and the spe-

cific integration of such antennas and the mm wave front ends associated with them depends on both form factor and cost constraints. This in turn impacts the coverage of the beams that can be synthesized by these antennas. Thus, while our simulations are for antenna arrays with isotropic elements for simplicity, the physical realization of the network node may impose constraints on network connectivity that must be taken into account. Another important topic for future research is the design of protocols that exploit the significant potential for spatial reuse enabled by the use of highly directional links, and for enabling co-existence of multiple WPANs in close proximity. Another interesting issue is detailed investigation of whether and how reflections can be used to steer around obstacles, as an alternative to, or in combination with, the use of relays. While the point of reflection can be thought of as a virtual relay, we note that the path followed by a reflection is constrained by the geometry of the environment, whereas the placement of an actual relay can be optimized (e.g., it can be put high up on a wall) to maximize connectivity. Much work also remains on detailed physical layer transceiver design for enabling beamsteering with minimal overhead and complexity, which includes cross-layer considerations both from below (RFIC and antenna design) and above (MAC design). Additionally, it is important to undertake design and performance evaluation with traffic models aimed specifically at some

of the applications driving the interest in high-speed WPANs, such as streaming compressed and uncompressed audio/video, as well as large file transfers.

A number of open problems need to be addressed to realize our vision of multiGigabit 60 GHz outdoor mesh networks. Joint cross-layer design of flow-control, routing, scheduling and MAC for highly directional mesh networks, including a comparison of centralized and distributed strategies, remains an important open issue. Detailed design, implementation, and testing of the envisioned omni-coverage yet highly directional mesh nodes is an important task, which involves design and validation of the antenna array patterns, low-complexity and low-overhead methods for electronic beamsteering, and low-power signal processing techniques for multiGigabit communication.

In the context of Sticky CSMA/CA, while we demonstrate large performance gains for a mix of intrinsically periodic traffic (such as VoIP) and data, further research is needed to extend Sticky CSMA/CA framework for efficient support of variable rate video over mesh networks. More broadly, the design of a comprehensive QoS solution for wireless mesh networks that employs Sticky CSMA/CA as a fundamental building block, with cross-layer interactions among application, network and MAC layers, is an interesting avenue for future research.

Many research questions remain on fully understanding the potential of learning and memory based algorithms for wireless protocol designs. For instance, can

we design stigmergy-inspired decentralized schemes that guarantee achieving a global network performance objective such as maximizing network utility for a given utility function? Can we ensure that our protocols react quickly enough to sudden changes in network traffic or topology (node or link failures)? Moreover, a clear understanding of the tradeoff between centralized options to resource management (which might be feasible for networks of moderate or small sizes) and our decentralized approaches remains to be understood.

For our resource biasing framework, important topics for future research include a detailed design and evaluation of decentralized protocols that implement our resource allocation strategies, leveraging the extensive recent work on protocol design for network utility maximization (NUM). It will be insightful to find out how much of the predicted gains for the mixed biasing strategies are actually realized via these protocols. Moreover, generalization of the analytical model to allow for multiple raw link speeds is another open issue, which may complicate the interaction between the local-scale and global-scale models.

Looking ahead, we believe that next generation wireless networks must be designed to meet the expectations of a large number of wireless users with increasingly savvy devices and diverse applications. Cross-layer wireless protocol designs are crucial to the eventual success of these networks: our work illustrates this ideology in the context of omnidirectional mesh networks operating in the 2.4

and 5.8 GHz WiFi bands as well as emerging highly directional 60 GHz wireless networks. In addition, our work also demonstrates the promise of implicit communication via tools of memory and learning in addressing challenging distributed coordination problems in wireless networks.

Bibliography

- [1] P. Gupta and P. R. Kumar, “The Capacity of Wireless Networks,” *IEEE Trans. Inform. Theory*, vol. 46, no. 2, pp. 388–404, 2000.
- [2] M. Chiang, S. Low, A. Calderbank, and J. Doyle, “Layering as Optimization Decomposition: A Mathematical Theory of Network Architectures,” *Proceedings of the IEEE*, vol. 95, no. 1, pp. 255–312, 2007.
- [3] X. Lin, N. B. Shroff, and R. Srikant, “A Tutorial on Cross-Layer Optimization in Wireless Networks,” *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 1452–1463, 2006.
- [4] IBMs 60-GHz Page. [Online]. Available: http://domino.watson.ibm.com/comm/research_projects.nsf/pages/mmwave.sixtygig.html
- [5] SiBEAM. [Online]. Available: <http://www.sibeam.com/>
- [6] S. Singh, F. Ziliotto, U. Madhow, E. Belding, and M. Rodwell, “Blockage and Directivity in 60 GHz Wireless Personal Area Networks: From Cross-Layer Model to Multihop MAC Design,” *IEEE J. Sel. Areas Commun.*, vol. 27, no. 8, pp. 1400–1413, Oct. 2009.
- [7] R. Mudumbai, S. Singh, and U. Madhow, “Medium Access Control for 60 GHz Outdoor Mesh Networks with Highly Directional Links,” in *Proc. IEEE INFOCOM 2009, Mini Conference*, April 2009, pp. 2871–2875.
- [8] S. Singh, R. Mudumbai, and U. Madhow, “Distributed Coordination with Deaf Neighbors: Efficient Medium Access for 60 GHz Mesh Networks,” *Submitted for publication*, 2009.
- [9] S. Singh, P. Acharya, U. Madhow, and E. Belding, “Sticky CSMA/CA: Implicit Synchronization and Real-time QoS in Mesh Networks,” *Ad Hoc Netw.*, vol. 5, no. 6, pp. 744–768, 2007.

- [10] S. Singh, U. Madhow, and E. Belding, “Beyond Proportional Fairness: A Resource Biasing Framework for Shaping Throughput Profiles in Multi-hop Wireless Networks,” in *Proc. IEEE INFOCOM 2008, Mini-Conference*, April 2008, pp. 2396–2404.
- [11] S. Singh, F. Ziliotto, U. Madhow, E. Belding, and M. Rodwell, “Millimeter Wave WPAN: Cross-Layer Modeling and Multi-Hop Architecture,” in *Proc. IEEE INFOCOM 2007, Mini Symposium*, May 2007, pp. 2336–2340.
- [12] S. Singh, U. Madhow, and E. M. Belding, “Shaping Throughput Profiles in Multihop Wireless Networks: A Resource Biasing Approach,” *Submitted for publication*, 2009.
- [13] P. Acharya, S. Singh, and H. Zheng, “Reliable Open Spectrum Communications through Proactive Spectrum Access,” in *Proc. First International Workshop on Technology and Policy for Accessing Spectrum*. ACM New York, NY, USA, 2006.
- [14] 60 GHz CMOS Radio Design at Berkeley Wireless Reaearch Center. [Online]. Available: http://bwrc.eecs.berkeley.edu/Research/RF/ogre_project/
- [15] WIGWAM - Wireless Gigabit with Advanced Multimedia Support. [Online]. Available: <http://www.wigwam-project.de/>
- [16] NICTA - Wireless GigE. [Online]. Available: http://www.nicta.com.au/research/projects/60ghz_wireless
- [17] N. Guo, R. C. Qiu, S. S. Mo, and K. Takahashi, “60-GHz Millimeter-Wave Radio: Principle, Technology, and New Results,” *EURASIP J. Wirel. Commun. Netw.*, vol. 2007, no. 1, pp. 48–48, 2007.
- [18] S. K. Yong and C.-C. Chong, “An Overview of Multigigabit Wireless through Millimeter Wave Technology: Potentials and Technical Challenges,” *EURASIP J. Wirel. Commun. Netw.*, vol. 2007, no. 1, pp. 1–10, 2007.
- [19] IEEE 802.15 WPAN Millimeter Wave Alternative PHY Task Group 3c (TG3c). [Online]. Available: <http://www.ieee802.org/15/pub/TG3c.html>
- [20] WirelessHD. [Online]. Available: <http://wirelesshd.org/>
- [21] ECMA International. [Online]. Available: <http://www.ecma-international.org/>

- [22] R. R. Choudhury, X. Yang, R. Ramanathan, and N. H. Vaidya, "On Designing MAC Protocols for Wireless Networks Using Directional Antennas," *IEEE Trans. Mob. Comput.*, vol. 5, no. 5, pp. 477–491, 2006.
- [23] R. Ramanathan, J. Redi, C. Santivanez, D. Wiggins, and S. Polit, "Ad Hoc Networking with Directional Antennas: A Complete System Solution," *IEEE J. Sel. Areas Commun.*, vol. 23, no. 3, pp. 496–506, March 2005.
- [24] T. Korakis, G. Jakllari, and L. Tassiulas, "CDR-MAC: A Protocol for Full Exploitation of Directional Antennas in Ad Hoc Wireless Networks," *IEEE Trans. Mob. Comput.*, vol. 7, no. 2, pp. 145–155, Feb. 2008.
- [25] A. Nasipuri, S. Ye, J. You, and R. Hiromoto, "A MAC Protocol for Mobile Ad Hoc Networks Using Directional Antennas," *Proc. IEEE WCNC 2000*, vol. 3, pp. 1214–1219 vol.3, 2000.
- [26] M. Takai, J. Martin, R. Bagrodia, and A. Ren, "Directional Virtual Carrier Sensing for Directional Antennas in Mobile Ad Hoc Networks," in *Proc. ACM MobiHoc 2002*. New York, NY, USA: ACM, 2002, pp. 183–193.
- [27] R. R. Choudhury and N. H. Vaidya, "Deafness: A MAC Problem in Ad Hoc Networks when Using Directional Antennas," *Proc. IEEE ICNP 2004*, pp. 283–292, 2004.
- [28] H. Singh and S. Singh, "Smart-Aloha for Multi-Hop Wireless Networks," *Mob. Netw. Appl.*, vol. 10, no. 5, pp. 651–662, 2005.
- [29] G. Jakllari, W. Luo, and S. V. Krishnamurthy, "An Integrated Neighbor Discovery and MAC Protocol for Ad Hoc Networks Using Directional Antennas," in *Proc. IEEE WOWMOM'05*, Washington, DC, 2005.
- [30] H. Xu, V. Kukshya, and T. S. Rappaport, "Spatial and Temporal Characteristics of 60 GHz Indoor Channels," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 3, pp. 620–630, Apr. 2002.
- [31] N. Moraitis and P. Constantinou, "Indoor Channel Measurements and Characterization at 60 GHz for Wireless Local Area Network Applications," *IEEE Trans. Antennas Propag.*, vol. 52, no. 12, pp. 3180–3189, Dec. 2004.
- [32] P. F. M. Smulders, "Broadband Wireless LANs: A Feasibility Study," Ph.D. dissertation, Eindhoven University of Technology, The Netherlands, 1995.

- [33] T. Manabe, K. Sato, H. Masuzawa, K. Taira, T. Ihara, Y. Kasashima, and K. Yamaki, "Effects of Antenna Directivity and Polarization on Indoor Multipath Propagation Characteristics at 60 GHz," *IEEE J. Sel. Areas Commun.*, vol. 14, no. 3, pp. 441–448, Apr. 1996.
- [34] H. Yang, M. Herben, and P. Smulders, "Impact of Antenna Pattern and Reflective Environment on 60 GHz Indoor Radio Channel Characteristics," *IEEE Antennas Wireless Propagat. Lett.*, vol. 4, pp. 300–303, Jun. 2005.
- [35] J. Kunisch, E. Zollinger, J. Pamp, and A. Winkelmann, "MEDIAN 60GHz Wideband Indoor Radio Channel Measurements and Model," in *Proc. IEEE VTC'99*, Amsterdam, The Netherlands, Sep. 1999.
- [36] T. Zwick, T. J. Beukema, and H. Nam, "Wideband Channel Sounder With Measurements and Model for the 60 GHz Indoor Radio Channel," *IEEE Trans. on Veh. Tech.*, vol. 54, no. 4, pp. 1266–1277, July 2005.
- [37] M.-S. Choi, G. Grosskopf, and D. Rohde, "Statistical Characteristics of 60 GHz Wideband Indoor Propagation Channel," *Proc. IEEE PIMRC 2005*, vol. 1, pp. 599–603, Sept. 2005.
- [38] C. Anderson and T. Rappaport, "In-Building Wideband Partition Loss Measurements at 2.5 and 60 GHz," *IEEE Tran. Wireless Commun.*, vol. 3, no. 3, pp. 922–928, May 2004.
- [39] A. Hammoudeh, D. Scammell, and M. Sanchez, "Measurements and Analysis of the Indoor Wideband Millimeter Wave Wireless Radio Channel and Frequency Diversity Characterization," *IEEE Trans. Antennas Propagat.*, vol. 51, no. 10, pp. 2974–2986, Oct. 2003.
- [40] P. F. M. Smulders, "Deterministic Modelling of Indoor Radio Propagation at 40-60 GHz," *Wireless Pers. Commun.*, vol. 1, no. 2, pp. 127–135, Jun. 1994.
- [41] F. Villanese, N. E. Evans, and W. G. Scanlon, "Pedestrian-Induced Fading for Indoor Channels at 2.45, 5.7 and 62GHz," in *Proc. 52th Vehicular Technology Conference, VTC'00*, Boston, MA, Sep. 2000.
- [42] M. Williamson, G. E. Athanasiadou, and A. R. Nix, "Investigating the Effects of Antenna Directivity on Wireless Indoor Communication at 60GHz," in *Proc. IEEE PIMRC'97*, Helsinki, Finland, Sep. 1997.

- [43] T. Manabe, K. Sato, H. Masuzawa, K. Taira, T. Ihara, Y. Kasashima, and K. Yamaki, "Polarization Dependence of Multipath Propagation and High-speed Transmission Characteristics of Indoor Millimeter-Wave Channel at 60 GHz," *IEEE Trans. Veh. Technol.*, vol. 44, no. 2, pp. 268–274, May 1995.
- [44] P. F. M. Smulders, "Exploiting the 60 GHz Band for Local Wireless Multimedia Access: Prospects and Future Directions," *IEEE Commun. Mag.*, vol. 40, no. 1, pp. 140–147, Jan. 2002.
- [45] S. Collonge, G. Zaharia, and G. Zein, "Influence of the Human Activity on Wide-Band Characteristics of the 60 GHz Indoor Radio Channel," *IEEE Tran. Wireless Commun.*, vol. 3, no. 6, pp. 2396–2406, Nov. 2004.
- [46] K. Sato and T. Manabe, "Estimation of Propagation-Path Visibility for Indoor Wireless LAN Systems under Shadowing Condition by Human Bodies," in *Proc. IEEE VTC'98*, Ottawa, Canada, May 1998.
- [47] R. Janaswamy, "An Indoor Pathloss Model at 60 GHz Based on Transport Theory," *IEEE Antennas Wireless Propagat. Lett.*, vol. 5, pp. 58–60, 2006.
- [48] S.-K. Yong, "TG3c Channel Modeling Sub-Committee Report," IEEE 802.15 WPAN Millimeter Wave Alternative PHY Task Group 3c (TG3c), Tech. Rep. IEEE 15-07-0584-01-003c, Mar. 2007.
- [49] J. D. Kraus, *Antennas for all Applications*. NY: McGraw-Hill, Inc., 2002.
- [50] U. Madhow, *Fundamentals of Digital Communication*. New York, NY: Cambridge University Press, 2008.
- [51] S. Bellofiore, J. Foutz, C. Balanis, and A. Spanias, "Smart-Antenna System for Mobile Communication Networks. Part 2. Beamforming and Network Throughput," *IEEE Antennas and Propagat. Mag.*, vol. 44, no. 4, pp. 106–114, Aug 2002.
- [52] A. Sayed, *Fundamentals of Adaptive Filtering*. IEEE Press, 2003.
- [53] S. J. Orfanidis, *Electromagnetic Waves and Antennas*. Piscataway, NJ: Online Book., Feb. 2008. [Online]. Available: <http://www.ece.rutgers.edu/~orfanidi/ewa/>
- [54] J. D. Kraus, *Electromagnetics*. New York: McGraw-Hill, Inc., 1991.
- [55] J. H. Whitteker, "Fresnel-Kirchhoff Theory Applied to Terrain Diffraction Problems," *Radio Science*, vol. 25, pp. 837–851, Oct. 1990.

- [56] J. Whitteker, "Physical Optics and Field-strength Predictions for Wireless Systems," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 3, pp. 515–522, Apr 2002.
- [57] T. S. Rappaport, *Wireless Communications: Principles and Practice*. Upper Saddle River, NJ: Prentice Hall, Inc., 2002.
- [58] J. D. Parsons, *The Mobile Radio Propagation Channel*. New York: John Wiley and Sons, Ltd., 2000.
- [59] T. Russell, C. Bostian, and T. Rappaport, "A Deterministic Approach to Predicting Microwave Diffraction by Buildings for Microcellular Systems," *IEEE Trans. Antennas Propagat.*, vol. 41, no. 12, pp. 1640–1649, Dec 1993.
- [60] C. Bettstetter, H. Hartenstein, and X. Pérez-Costa, "Stochastic Properties of the Random Waypoint Mobility Model," *Wirel. Netw.*, vol. 10, no. 5, pp. 555–567, 2004.
- [61] QualNet Network Simulator. [Online]. Available: <http://www.scalable-networks.com>
- [62] "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) specifications Amendment 8: Medium Access Control (MAC) Quality of Service Enhancements," IEEE Std. 802.11e-2005, November 2005.
- [63] S. Alalusi and R. Brodersen, "A 60GHz Phased Array in CMOS," in *Proc. IEEE CICC '06*, Sept. 2006, pp. 393–396.
- [64] D. Liu and R. Sirdeshmukh, "A Patch Array Antenna for 60 GHz Package Applications," in *Proc. IEEE AP-S Symposium '08*, July 2008, pp. 1–4.
- [65] Y.-B. Ko, V. Shankarkumar, and N. Vaidya, "Medium Access Control Protocols Using Directional Antennas in Ad Hoc Networks," *Proc. IEEE INFOCOM 2000*, vol. 1, pp. 13–21 vol.1, 2000.
- [66] S. Kulkarni and C. Rosenberg, "DBSMA: A MAC Protocol for Multi-Hop Ad-Hoc Networks with Directional Antennas," in *Proc. IEEE PIMRC'05*, vol. 2, Sep 2005, pp. 1371–1377.
- [67] Z. Huang, C.-C. Shen, C. Srisathapornphat, and C. Jaikaeo, "A Busy-Tone Based Directional MAC Protocol for Ad Hoc Networks," *Proc. MILCOM 2002.*, vol. 2, pp. 1233–1238 vol.2, Oct. 2002.

- [68] S. Sanghavi, L. Bui, and R. Srikant, “Distributed Link Scheduling with Constant Overhead,” in *Proc. ACM SIGMETRICS*, 2007, pp. 313–324.
- [69] W. Wang, Y. Wang, X.-Y. Li, W.-Z. Song, and O. Frieder, “Efficient Interference-Aware TDMA Link Scheduling for Static Wireless Networks,” in *Proc. MobiCom '06*. NY, USA: ACM, 2006, pp. 262–273.
- [70] T. Salonidis and L. Tassiulas, “Distributed Dynamic Scheduling for End-to-End Rate Guarantees in Wireless Ad Hoc Networks,” in *Proc. ACM MobiHoc'05*. New York, NY, USA: ACM, 2005, pp. 145–156.
- [71] Y. Yi, G. de Veciana, and S. Shakkottai, “Learning Contention Patterns and Adapting to Load/Topology Changes in a MAC Scheduling Algorithm,” in *Proc. IEEE WiMesh '06*, Reston, VA, USA, 2006.
- [72] R. Solis, V. Borkar, and P. Kumar, “A New Distributed Time Synchronization Protocol for Multihop Wireless Networks,” in *Proc. 45th IEEE CDC, San Diego, CA, USA*, 2006.
- [73] P. Sommer and R. Wattenhofer, “Gradient Clock Synchronization in Wireless Sensor Networks,” in *Proc. ACM/IEEE IPSN '09*, 2009.
- [74] R. Mudumbai and U. Madhoo, “Information Theoretic Bounds for Sensor Network Localization,” in *Proc. IEEE ISIT'08*, July 2008, pp. 1602–1606.
- [75] J. Zander, “Slotted ALOHA Multihop Packet Radio Networks with Directional Antennas,” *Electron. Lett.*, vol. 26, no. 25, pp. 2098–2100, Dec. 1990.
- [76] J. Ward and J. Compton, R.T., “Improving the Performance of a Slotted ALOHA Packet Radio Network with an Adaptive Array,” *Communications, IEEE Transactions on*, vol. 40, no. 2, pp. 292–300, Feb 1992.
- [77] H. N. Gabow, “Implementation of Algorithms for Maximum Matching on Nonbipartite Graphs,” Ph.D. dissertation, Stanford University, Stanford, CA, USA, 1974.
- [78] M. Leconte, J. Ni, and R. Srikant, “Improved Bounds on the Throughput Efficiency of Greedy Maximal Scheduling in Wireless Networks,” in *Proc. ACM MobiHoc '09*. New York, NY, USA: ACM, 2009, pp. 165–174.
- [79] R. Jain, D. Chiu, and W. Hawe, “A Quantitative Measure of Fairness and Discrimination for Resource Allocation in Shared Systems,” Digital Equipment Corp, DEC-TR-301, Tech. Rep., 1984.

- [80] Y. Sun, I. Sheriff, E. M. Belding-Royer, and K. C. Almeroth, “An Experimental Study of Multimedia Traffic Performance in Mesh Networks,” in *WiTMeMo’05: Workshop on Wireless Traffic Measurements and Modeling*, Seattle, WA, June 2005.
- [81] F. Kelly, “Modelling Communication Networks, Present and Future,” *Philosophical Transactions of the Royal Society of London*, vol. A354, no. 1707, pp. 437–463, 1996.
- [82] T. Henderson, D. Kotz, and I. Abyzov, “The Changing Usage of a Mature Campus-wide Wireless Network,” in *Proc. 10th Annual International Conference on Mobile Computing and Networking (MobiCom’04)*, Philadelphia, PA, September 2004.
- [83] Cisco 7900 Series IP Phones. [Online]. Available: <http://www.cisco.com/en/US/products/hw/phones/ps379/>
- [84] “Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications,” IEEE Std. 802.11, January 1997.
- [85] J. L. Sobrinho and A. S. Krishnakumar, “Real-time Traffic over the IEEE 802.11 Medium Access Control Layer,” *Bell Labs Technical Journal*, vol. 1, no. 2, pp. 172–187, Autumn 1996.
- [86] N. H. Vaidya, P. Bahl, and S. Gupta, “Distributed Fair Scheduling in a Wireless LAN,” in *Proc. 6th Annual International Conference on Mobile Computing and Networking (MobiCom’00)*, Boston, MA, August 2000.
- [87] S. J. Golestani, “A Self-Clocked Fair Queueing Scheme for Broadband Applications,” in *Proc. IEEE INFOCOM’94*, Toronto, Canada, June 1994.
- [88] C. R. Lin and M. Gerla, “Asynchronous Multimedia Multihop Wireless Networks,” in *Proc. IEEE INFOCOM’97*, Kobe, Japan, April 1997.
- [89] A. Rao and I. Stoica, “An Overlay MAC Layer for 802.11 Networks,” in *MobiSys ’05: Proc. of the 3rd International Conference on Mobile Systems, Applications, and Services*, Seattle, WA, 2005.
- [90] I. Rhee, A. Warriar, M. Aia, and J. Min, “Z-MAC: A Hybrid MAC for Wireless Sensor Networks,” in *SenSys ’05: Proc. of the 3rd International Conference on Embedded Networked Sensor Systems*, San Diego, CA, 2005.

- [91] C.-H. Lin, H. Dong, U. Madhow, and A. Gersho, "Supporting Real-time Speech on Wireless Ad hoc Networks: Inter-packet Redundancy, Path diversity, and Multiple Description Coding," in *Proc. 2nd ACM International Workshop on Wireless Mobile Applications and Services on WLAN Hotspots (WMASH'04)*, Philadelphia, PA, October 2004.
- [92] A. Vasan, R. Ramjee, and T. Woo, "ECHOS -Enhanced Capacity 802.11 Hotspots," in *Proc. IEEE INFOCOM'05*, Miami, FL, March 2005.
- [93] "Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications: Higher Speed Physical Layer (PHY) Extension in the 2.4GHz Band," IEEE Std. 802.11b/D5.0, April 1999.
- [94] K. Sanzgiri, I. D. Chakeres, and E. M. Belding-Royer, "Determining Intra-Flow Contention along Multihop Paths in Wireless Networks," in *Proc. 1st Annual International Conference on Broadband Networks (BROADNETS'04)*, San Jose, CA, October 2004.
- [95] G. Bianchi, "Performance Analysis of the IEEE 802.11 Distributed Coordination Function," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 3, pp. 535–547, March 2000.
- [96] Y. C. Tay and K. C. Chua, "A Capacity Analysis for the IEEE 802.11 MAC Protocol," *Wireless Networks*, vol. 7, no. 2, pp. 159–171, March 2001.
- [97] H. S. Chhaya and S. Gupta, "Performance Modeling of Asynchronous Data-transfer Methods of IEEE 802.11 MAC Protocol," *Wireless Networks*, vol. 3, no. 3, pp. 217–234, August 1997.
- [98] M. M. Carvalho and J. J. Garcia-Luna-Aceves, "A Scalable Model for Channel Access Protocols in Multihop Ad hoc Networks," in *Proc. 10th Annual International Conference on Mobile Computing and Networking (MobiCom'04)*, Philadelphia, PA, September 2004.
- [99] MeshDynamics—High Performance Metro Mesh Networks with Structured Mesh (TM). [Online]. Available: <http://www.meshdynamics.com>
- [100] A. P. Markopoulou, F. A. Tobagi, and M. J. Karam, "Assessing the Quality of Voice Communications over Internet Backbones," *IEEE/ACM Trans. Netw.*, vol. 11, no. 5, pp. 747–760, October 2003.
- [101] R. Zopf, "Real-time Transport Protocol (RTP) Payload for Comfort Noise (CN)," RFC 3389, Sep. 2002.

- [102] P. T. Brady, "A Model for Generating On-Off Speech Patterns in Two-way Conversation," *The Bell System Technical Journal*, vol. 48, no. 9, pp. 2445–2472, Sep. 1969.
- [103] S. Garg and M. Kappes, "Can I Add a VoIP Call?" in *Proc. IEEE International Conference on Communications (ICC 2003)*, Anchorage, AL, May 2003.
- [104] —, "An Experimental Study of Throughput for UDP and VoIP Traffic in IEEE 802.11b Networks," in *IEEE Wireless Communications and Networking Conference (WCNC 2003)*, New Orleans, LA, 2003.
- [105] M. Takai, J. Martin, and R. Bagrodia, "Effects of Wireless Physical Layer Modeling in Mobile Ad hoc Networks," in *Proc. ACM MobiHoc 2001*, Long Beach, CA, October 2001.
- [106] H.-K. Wu, C.-H. Hung, M. Gerla, and R. Bagrodia, "Speech Support in Wireless, Multihop Networks," in *ISPAN '97: Proc. 1997 International Symposium on Parallel Architectures, Algorithms and Networks*. Washington, DC, USA: IEEE Computer Society, 1997.
- [107] M. Grossglauser, S. Keshav, and D. N. C. Tse, "RCBR: A Simple and Efficient Service for Multiple Time-Scale Traffic," *IEEE/ACM Trans. Netw.*, vol. 5, no. 6, pp. 741–755, 1997.
- [108] F. Kelly, A. Maulloo, and D. Tan, "Rate Control In Communication Networks: Shadow Prices, Proportional Fairness and Stability," *J. Operational Research Soc.*, vol. 49, pp. 237–252, 1998.
- [109] J. Mo and J. Walrand, "Fair End-to-End Window-based Congestion Control," *IEEE/ACM Trans. Netw.*, vol. 8, no. 5, pp. 556–567, 2000.
- [110] L. Massouli and J. Roberts, "Bandwidth Sharing: Objectives and Algorithms," *IEEE/ACM Trans. Netw.*, vol. 10, no. 3, pp. 320–328, 2002.
- [111] L. Tassiulas and S. Sarkar, "Max-Min Fair Scheduling in Wireless Networks," in *Proc. IEEE INFOCOM'02*, New York City, NY, June 2002, pp. 763–772.
- [112] U. Akyol, M. Andrews, P. Gupta, J. D. Hobby, I. Saniee, and A. L. Stolyar, "Joint Scheduling and Congestion Control in Mobile Ad-Hoc Networks," in *Proc. IEEE INFOCOM 2008*, 2008, pp. 619–627.

- [113] A. Eryilmaz and R. Srikant, “Joint Congestion Control, Routing, and MAC for Stability and Fairness in Wireless Networks,” *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 1514–1524, 2006.
- [114] L. Bui, R. Srikant, and A. Stolyar, “Novel Architectures and Algorithms for Delay Reduction in Back-pressure Scheduling and Routing,” in *Proc. IEEE INFOCOM 2008 Mini Conference*, 2009.
- [115] M. Neely, E. Modiano, and C.-P. Li, “Fairness and Optimal Stochastic Control for Heterogeneous Networks,” *IEEE/ACM Trans. Netw.*, vol. 16, no. 2, pp. 396–409, April 2008.
- [116] J. Lee, M. Chiang, and R. A. Calderbank, “Jointly Optimal Congestion and Contention Control Based on Network Utility Maximization,” *IEEE Commun. Lett.*, vol. 10, no. 3, pp. 216–218, 2006.
- [117] X. Lin and N. Shroff, “The Impact of Imperfect Scheduling on Cross-Layer Congestion Control in Wireless Networks,” *IEEE/ACM Trans. Netw.*, vol. 14, no. 2, pp. 302–315, 2006.
- [118] L. Chen, S. H. Low, and J. C. Doyle, “Joint Congestion Control and Media Access Control Design for Ad Hoc Wireless Networks,” in *Proc. IEEE INFOCOM 2005.*, vol. 3, 2005.
- [119] Y. Yi and S. Shakkottai, “Hop-by-Hop Congestion Control over a Wireless Multi-Hop Network,” *IEEE/ACM Trans. Netw.*, vol. 15, no. 1, pp. 133–144, 2007.
- [120] M. Chiang, “Balancing Transport and Physical Layers in Wireless Multihop Networks: Jointly Optimal Congestion Control and Power Control,” *IEEE J. Sel. Areas Commun.*, vol. 23, no. 1, pp. 104–116, 2005.
- [121] R. Srikant, *The Mathematics of Internet Congestion Control*. Birkhauser, 2004.
- [122] L. Tassiulas and A. Ephremides, “Stability Properties of Constrained Queueing Systems and Scheduling Policies for Maximum Throughput in Multihop Radio Networks,” *IEEE Trans. Autom. Control*, vol. 37, no. 12, pp. 1936–1948, 1992.
- [123] B. Radunovic and J.-Y. L. Boudec, “Rate Performance Objectives of Multihop Wireless Networks,” *IEEE Trans. Mob. Comput.*, vol. 3, no. 4, pp. 334–349, 2004.

- [124] A. Velayutham, K. Sundaresan, and R. Sivakumar, “Non-Pipelined Relay Improves Throughput Performance of Wireless Ad-Hoc Networks.” in *Proc. IEEE INFOCOM’05*, Miami, FL, March 2005, pp. 477–490.
- [125] N. Zhou, H. Wu, and A. A. Abouzeid, “The Impact of Traffic Patterns on the Overhead of Reactive Routing Protocols,” *IEEE J. Sel. Areas Commun.*, vol. 23, no. 3, pp. 547–560, Mar. 2005.
- [126] J. Li, C. Blake, D. S. J. De Couto, H. I. Lee, and R. Morris, “Capacity of Ad Hoc Wireless Networks,” in *Proc. ACM Mobicom 2001*, Rome, Italy, July 2001, pp. 61–69.
- [127] M. Heusse, F. Rousseau, G. Berger-Sabbatel, and A. Duda, “Performance Anomaly of 802.11b,” in *Proc. IEEE INFOCOM’03*, San Francisco, USA, March-April 2003.
- [128] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge university press, 2004.
- [129] A. Stolyar, “Maximizing Queueing Network Utility Subject to Stability: Greedy Primal-Dual Algorithm,” *Queueing Systems*, vol. 50, no. 4, pp. 401–457, 2005.